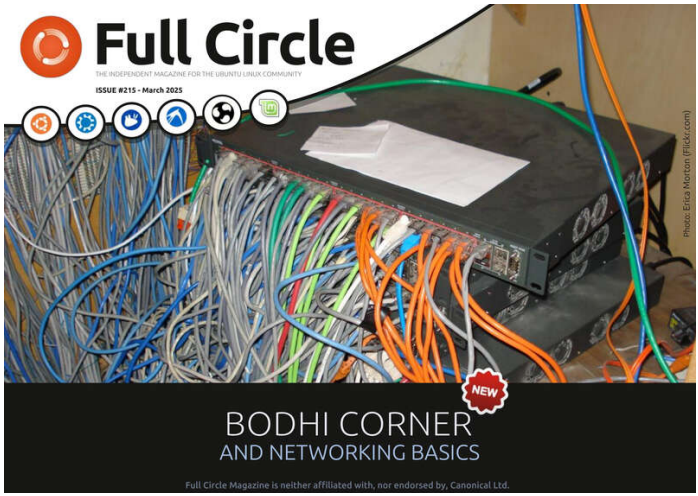
Full Circle
THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY
ISSUE #215 - March 2025
Photo: Erica Morton (Flickr.com)
NEW
BODHI CORNER
AND NETWORKING BASICS
Full Circle Magazine is neither affiliated with, nor endorsed by, Canonical Ltd.

Issue #215

Table of Contents

1. Editorial
2. News
3. Command & Conquer:
4. HowTo: Network Basics
5. HowTo: Trading up to Linux – Part 3
6. HowTo: Latex - Babel
7. Kilobyte magazine advert
8. HowTo: Drawing with Inkscape - Part 155
9. shells.com advert
10. BODHI CORNER
11. UBports Devices: News From The Community
12. InEasySteps advert
13. My Opinion: To Rust or not to Rust
14. nostarch.com advert
15. HowTo: Write For Full Circle Magazine
16. Review: DistroSea
17. Review: VirtualBox vs Gnome Boxes
18. Letters
19. Q&A
20. Ubuntu Games: Fates of Ort
21. Patrons
22. How to Contribute

Landmarks

1. *Table of Contents*

The articles contained in this magazine are released under the Creative Commons Attribution-Share Alike 3.0 Unported license. This means you can adapt, copy, distribute and transmit the articles but only under the following conditions: you must attribute the work to the original author in some way (at least a name, email or URL) and to this magazine by name ('Full Circle Magazine') and the URL *www.fullcirclemagazine.org* (but not attribute the article(s) in any way that suggests that they endorse you or your use of the work). If you alter, transform, or build upon this work, you must distribute the resulting work under the same, similar or a compatible license. Full Circle magazine is entirely independent of Canonical, the sponsor of the Ubuntu projects, and the views and opinions in the magazine should in no way be assumed to have Canonical endorsement.

Please note: articles in this magazine are provided with absolutely no warranty whatsoever; neither the contributors nor Full Circle Magazine accept any responsibility or liability for loss or damage resulting from readers choosing to apply this content to theirs or others computers and equipment.

RETURN TO CONTENTS

Editorial

Welcome to the latest issue of Full Circle

This month, we bring you more Latex, Trading Up, and Inkscape. Learn About will return next issue. In its place we have a piece (also from Erik) on networking basics. And, yes, articles are still up there with hen's teeth. So if you have any articles you can spare, please send them to: *ronnie@fullcirclemagazine.org*.

No reviews from Adam this month. You know what that means, right? Yep. We'll soon be on to the next version. Next issue will probably start the 25.04 reviews. We do still have a couple of reviews for you, though; DistroSea and VirtualBox vs Gnome Boxes. And, of course, a game review.

Full Circle Weekly News presenter extraordinaire Moss Bliss is starting a Bodhi Corner this month. If you're curious about the Bodhi distro, then that's the place to be.

Remember: the Full Circle Weekly News is available on Spotify and YouTube. The more upvotes and reviews you give it on those platforms the more exposure we get. And, we have a Table of Contents which lists every article from every issue of FCM. Huge thanks to Paul Romano for maintaining: *https://goo.gl/tpOKqm* and, if you're looking for some help, advice, or just a chinwag: remember that we have a Telegram group: *https://t.me/joinchat/24ec1oMFO1ZjZDc0*. I hope to see you there. Come and say hello.

All the best for 2025! Ronnie
ronnie@fullcirclemagazine.org

FCM PATREON : https://www.patreon.com/fullcirclemagazine

RETURN TO CONTENTS

NEWS

Submitted by ErikTheUnread

Full Circle Weekly News

http://fullcirclemagazine.org/feed/podcast

http://www.stitcher.com/s?fid=85347&refid=stpr

http://tunein.com/radio/Full-Circle-Weekly-News-p855064/

https://player.fm/series/the-full-circle-weekly-news

Week 9

Lutris 0.5.19 Released:

24/02/2025

The Lutris 0.5.19 release of the gaming platform has been published. It provides tools to simplify the installation, configuration and management of games on Linux. The project code is written in Python and is distributed under the GPLv3 license. Ready-made builds are available only in the flatpak format :(and Ubuntu is five (5) versions behind as usual! (just checked Synaptic, yep, 5.14)

The project maintains a catalog for quick search and installation of gaming applications, allowing one click to launching games in Linux, through a single interface, without worrying about installing dependencies and settings. Runtime components for launching games are supplied by the project and are not tied to the distribution used. Runtime is a distribution-independent set of libraries, including components from SteamOS and Ubuntu, as well as various additional libraries.

You can install games distributed through GOG, Steam, Epic Games Store, Battle.net, Amazon Games, Origin and Uplay. However, Lutris itself only acts as an intermediary and does not sell games, so for commercial games the user must independently purchase the game from the corresponding service (free games can be launched with one click from the Lutris graphical interface).

Each game in Lutris is tied to a boot script and a handler that describes the environment for launching the game. This includes ready-made profiles with optimal settings for launching games under Wine. In addition to Wine, games can be launched using game console emulators such as RetroArch, Dosbox, FS-UAE, ScummVM, MESS/MAME and Dolphin .

https://github.com/lutris/lutris/releases/tag/v0.5.19

Apache NetBeans 25 IDE:

24/02/2025

The release of Apache NetBeans 25 integrated development environment is presented, which provides support for programming languages, like Java SE, Java EE, PHP, C/C++, JavaScript, Rust and Groovy. Ready-made builds are generated for Linux (snap , flatpak), Windows and macOS.

https://netbeans.apache.org/front/main/blogs/entry/announce-apache-netbeans-25-released/

Release of MythTV 35:

24/02/2025

After a year of development, the MythTV 35 home media center platform was released, allowing you to turn your desktop PC into a TV, video recording system, music center, photo album, DVD recording and viewing station. The project code is written in C++ and is distributed under the GPLv2 license.

The MythTV architecture is based on the division of the backend for storing or capturing video (IPTV, DVB cards, etc.) and the frontend for generating the interface (web interface and Qt-based GUI). The frontend is capable of working simultaneously with several backends, which can be launched both on the local system and on external computers. Extended functionality is implemented through plugins. Currently, two sets of plugins are available - official and unofficial. The range of capabilities covered by the plugins ranges from integration with various online services to tools for working with a web camera and organizing video communication between PCs.

https://www.mythtv.org/news/175/v35.0%2520Released

Aqualung 2.0 Music Player Ported to GTK3 Released:

25/02/2025

The release of the Aqualung 2.0 music player has been published. It provides an interface for playing audio CDs, local music collection files, online radio stations and podcasts. Aqualung's special feature is its gapless playback mode, which allows playing music files in a continuous stream, in which one composition replaces another without pausing. The project code is written in C and is distributed under the GPLv2 license. It supports Linux, FreeBSD, OpenBSD, Windows and macOS.

Among the advanced features, the following was highlighted, the ability to increase or decrease the sampling frequency of the file being played before sending the audio stream to the audio output device. LADSPA plug-ins can be connected to apply effects, improve the quality and process the sound . For example, you can use plug-ins with different implementations of equalizers, spatial sound optimizers and tube preamplifier simulators.

The key change in Aqualung 2.0 was the transition of the application from GTK2 to the GTK3 branch. CSS was used to design the interface. The ability to use previously prepared skins to change the interface was discontinued. An option to switch between light and dark themes was added to the settings. Support for new versions of the Monkey's Audio format for encoding sound without loss of quality was added. The number of supported formats in the FFmpeg-based decoder was expanded.

https://github.com/jeremyevans/aqualung/releases/tag/2.0

X.Org Server 21.1.16 Update Fixes 8 Vulnerabilities:

26/02/2025

Corrective releases of X.Org Server 21.1.16 and the DDX (Device-Dependent X) component xwayland 24.1.6 have been published, ensuring the launch of X.Org Server for the execution of X11 applications in Wayland-based environments. The new version of X.Org Server fixes 8 vulnerabilities. The

problems can potentially be exploited for privilege escalation on systems where the X server is executed with root rights, as well as for remote code execution in configurations where X11 session redirection using SSH is used for access.

https://gitlab.freedesktop.org/xorg/xserver/-/tags/

Nextcloud Hub 10 is released:

26/02/2025

The Nextcloud Hub 10 platform is released, providing a self-sufficient solution for organizing the collaboration of employees of enterprises and teams developing various projects. At the same time, the Nextcloud 31 cloud platform underlying Nextcloud Hub was published, allowing you to deploy cloud storage with support for synchronization and data exchange, providing the ability to view and edit data from any device at any point in the network (download a list of participants in CSV format, create, import and export surveys using a web interface or WebDAV). The Nextcloud server can be deployed on any hosting that supports the execution of PHP scripts and provides access to SQLite, MariaDB / MySQL or PostgreSQL. The Nextcloud source code is distributed under the AGPL license.

In terms of the tasks it solves, Nextcloud Hub resembles Google Docs and Microsoft 365, but allows you to deploy a fully controlled collaborative infrastructure that operates on its own servers and is not tied to external cloud services. Nextcloud Hub combines several open -source add-on applications over the Nextcloud cloud platform into a single environment, allowing you to work together with office documents, files, and information for planning tasks and events. The platform also includes add-ons for accessing email, exchanging messages, organizing video conferences, and chats.

User authentication can be performed either by a local database or through integration with LDAP/Active Directory, Kerberos, IMAP and Shibboleth/SAML 2.0, including the use of two-factor authentication, SSO (Single-sign-on) and linking

new systems to an account using a QR code. Version control allows you to track changes in files, comments, sharing rules and tags.

https://github.com/nextcloud/server/releases/tag/v31.0.0

A Project to Run FreeBSD Programs on Linux:

27/02/2025

FreeBSD developers have published a development report for the fourth quarter of 2024, which mentions the bsd-user-4-linux project, which is developing tools for running applications built for FreeBSD on Linux. The goal of the project is to achieve the ability to build natively using FreeBSD tools on Linux, as well as the ability to build packages for FreeBSD on Linux using the native FreeBSD build tools.

To run FreeBSD executable files, a fork of the QEMU emulator is used, operating in 'User Mode' Emulation mode, designed to run processes compiled for other CPUs and operating systems. In this mode, QEMU translates system calls and processes signals. To run applications, you need to deploy libraries and settings from the FreeBSD base system to a local directory. The project can be considered as BSDlator - a reverse of Linuxulator.

At the current stage of development, the launch of the main system utilities (sh, bash, find, grep, git, clang, etc.) works, dynamic linking and shared libraries are supported, network functions are available. For example, you can already rebuild FreeBSD with the command "make -j80 buildworld" while in Linux. The missing functionality includes the inability to launch the GDB debugger, the unavailability of IPC, kevent/kqueue functions, and some sysctl (for example, name2oid).

Additionally, the project is preparing container images with the FreeBSD working environment for various architectures (linux/386, linux/amd64, linux/arm/v5 and linux/arm64/v8) and providing GitHub Actions for creating such images.

https://www.freebsd.org/status/report-2024-10-2024-12/

OpenCloud 1.0 Now Available:

27/02/2025

The first release of the OpenCloud platform has been presented, allowing you to deploy a system for file sharing and collaborative work on content on your server. The project is presented as an open alternative to proprietary systems Microsoft SharePoint, Google Drive and Dropbox, compliant with GDPR requirements.

OpenCloud is a fork of the OCIS (ownCloud Infinite Scale) platform, rewritten from PHP to Go, unlike the original ownCloud codebase and the NextCloud project that branched off from it. The creators of OpenCloud tried to rid the codebase of unnecessary functionality and focus on the highest quality implementation of the main task - collaborative work with files.

In addition to file storage and sharing features, as well as synchronized access to a file collection from different devices, OpenCloud includes capabilities for collaborative editing of documents in real time, integration with the Collabora Online office suite and Markdown Editor (ToastUI), text extraction from images and scanned documents using OCR, and the use of the ICAP protocol to check downloaded files in antivirus packages.

You can limit the lifetime of published files, provide access via a link, and protect the content with a password. Tags, filters, and full-text search are supported to simplify file navigation. The system tracks the entire history of work with files and provides support for rolling back changes to a specific version in the past. You can organize teams by assigning subdirectories ("workspaces") to individual teams.

The administrator is provided with a detailed interface for managing access rights and users (for example, you can allow a user to only view or only upload to storage). Role-Based Access

Control (RBAC) is supported. The platform uses the Privacy-First architecture, which means that an OpenCloud user with administrator rights cannot access user content.

The server part is written in Go, distributed under the Apache 2.0 license and supports the WebDAV, gRPC, Microsofts RESTful Web API Graph, OCS, OCM 1.1 and OpenID Connect programming interfaces. It is noted that the server is implemented using the microservices concept and can be scaled from installation on Raspberry Pi boards to large multi-server implementations.

The desktop client is written in C++ using Qt, published under the GPLv3 license, and supports building for Windows, macOS, and Linux. The client also has a built-in file synchronization feature and mounting of shared storage as a virtual file system.

https://opencloud.eu/en/news/opencloud-now-available-new-open-source-alternative-microsoft-sharepoint

Linux kernel 6.15 to remove SystemV filesystem:

27/02/2025

SUSE developer Jan Kara has proposed a patch that removes code from the Linux kernel that supports Xenix, SystemV/386, and Coherent file systems. The code has been orphaned since 2023, and no one has been willing to maintain it since then. Kernel maintainer Christian Brauner, who is responsible for file systems, agreed with the proposal and accepted the patch into the vfs-6.15 branch. A few hours ago, the contents of vfs-6.15 were moved to the linux-next branch, which is preparing changes for the upcoming 6.15 kernel.

The story begins in 2002, when the Linux kernel got rid of the global lock (BKL - Big Kernel Lock) in favor of more granular locks. At that time, the SystemV file system switched to rwlock, but this led to problems due to which the processor could not go to sleep during the call of the entire function that read the inode metadata. The problem was noticed only in 2023 as a

result of stress testing with the SyzBot tool. Such an incident shows that no one uses the SystemV file system and it can be removed.

https://lore.kernel.org/all/20250221-dienlich-metapher-c2755e73b3f7@brauner/T/%23mf04463bbb9bb2d9461ea135955cd33ac29ebc0cd

Electronic Arts Opens Command & Conquer Games Under GPL License:

28/02/2025

Electronic Arts has announced the full source code for the Command & Conquer series of games: Tiberian Dawn, Red Alert, Renegade, Generals, and Generals Zero Hour. Tiberian Dawn, Red Alert, and Generals are real-time strategy games that feature military confrontations in various alternate realities and require building up military power by extracting resources available on the map. Renegade is a 3D first-person shooter with strategy elements.

The release of the code was made possible by the work of Luke Feenan, an administrator of the CnCNet community that preserves the legacy of the Command & Conquer games. Electronic Arts gave Luke access to their archives, and he did the work of restoring the state of the code from the Perforce repositories and getting it into a buildable state.

The build is supported only for Windows (the DirectX graphics API is used). To build C&C Renegade and C&C Generals, the Microsoft Visual Studio C++ 6+ compiler is required. The C&C Tiberian Dawn and C&C Red Alert code is tied to the Watcom C/C++ compiler (v10.6) and requires additional work to recreate the working build environment. Game resources are not included in the kit, so to run games based on executable files built from the code, resources from the original games are required.

The code is open sourced under the GPLv3 license with additional terms regarding Electronic Arts trademarks.

Trademark rights are not transferred with the code, and the license does not grant rights to distribute modified versions of the programs using Electronic Arts trademarks, such as "Command & Conquer". Products created based on the published code must not state affiliation with Electronic Arts or its employees, and must not misrepresent the history of their origin. When creating modifications, the name of the games must be changed, and a note must mention that the product is a modified version of Electronic Arts code.

Opening up the game under the GPLv3 license allows using components from the original Command & Conquer games to extend the functionality of the existing open source projects CnCNet and Open RA , which develop an unofficial network stack and open source engine for C&C Tiberian Dawn, C&C Red Alert, and Dune 2000. The game requires the original game assets to run, and can be used in accordance with the mod developer agreement, which allows non-commercial use only.

https://www.ea.com/games/command-and-conquer/command-and-conquer-remastered/news/steam-workshop-support

Fish 4.0 Shell Rewritten in Rust:

27/02/2025

The interactive command shell, fish 4.0.0 (friendly interactive shell) has been released. It is aimed at creating a more user-friendly alternative to bash and zsh. Fish supports features such as syntax highlighting with automatic detection of input errors, suggestion of possible input options based on the history of past operations, autocompletion of input options and commands using their description in man-guides, comfortable work out of the box without the need for additional configuration, a simplified scripting language, clipboard support, and search tools in the history of completed operations. The project code is distributed under the GPLv2 license. Ready-made packages are formed for Ubuntu, Debian, Fedora and openSUSE.

The new branch is notable for the complete translation of the

code base from C++ to Rust. It is noted that the transition to Rust allowed solving problems with multithreading, increasing the safety of working with memory, making the project more attractive to new developers, and using modern tools that detect errors at the compilation stage. The work on rewriting fish in Rust took two years. More than 200 developers took part in the creation of the 4.0 release.

https://fishshell.com/blog/new-in-40/

NetworkManager 1.52.0 Released:

01/03/2025

A stable release of the interface for simplifying network settings, has been published - NetworkManager 1.52.0. Plugins for VPN support (Libreswan, OpenConnect, Openswan, SSTP, etc.) are being developed within their own development cycles. (the post is basically just a list of stuff, like" Added support for using FEC (Forward Error Correction) mode when calling the "ethtool" utility. Link included below.)

https://gitlab.freedesktop.org/NetworkManager/NetworkManager/-/tags/1.52.0

Release of Brushshe 1.2.0:

02/03/2025

The release of the lightweight raster graphics editor, Brushshe 1.2.0, has been published Features include stickers, frames and effects. The program supports saving drawings in different formats and has its own gallery where you can view saved drawings. The project code is written in Python using the CustomTkinter graphical toolkit and is distributed under the GNU GPL v3 license. There is a build for Windows, tested in Windows 11, Windows 10 and Wine (Ubuntu users can run it via Python).

https://github.com/limafresh/Brushshe/releases/tag/v1.2.0

GhostBSD Release 25.01:

28/02/2025

The release of the desktop-oriented GhostBSD 25.01 distribution is presented. It is based on FreeBSD 14 and offers builds with MATE and Xfce user environments. By default, GhostBSD uses the ZFS file system. Both Live mode and installation on a hard drive are supported (using the proprietary ginstall installer written in Python). Boot images are generated for the x86_64 architecture (2.8 GB with MATE and 2.6 GB with Xfce).

https://ghostbsd.org/news/GhostBSD_25.01-R14.2p1_Now_Available

ROSA Fresh 13 distribution is available:

28/02/2025

Three years after the formation of the previous major branch, the company NTC IT ROSA published the ROSA Fresh 13 distribution, based on the new rosa 13 platform. ROSA Fresh 13 is positioned as the first release demonstrating the capabilities of the new platform and intended primarily for enthusiasts. Builds with KDE (3 GB) and GNOME (3 GB) desktops for the x86_64 architecture are available for download. The repository contains packages for the aarch64, e2kv4, i686, loongarch64, riscv64 and x86_64 architectures.

The new release is notable for updating packages to fresh versions of programs. The most significant change was the transition of the desktop environment to KDE 6 components (KDE Plasma 6.3.1 release is used). At the same time, packages with KDE Plasma 5.27.12 were left in the distro as an option. GNOME 47, LXQt 2.1, Xfce 4.20, MATE 1.28 were updated. The distribution uses the releases of Mesa 24.3, Qt 6.8.2, GCC 14.2, LLVM 19, Glibc 2.40, systemd 256, Perl 5.38, Python 3.11, Rust 1.84, Ruby 3.2.2, Go 1.23.6, PHP 8.3.14, Libreoffice 24.8.5, gimp 3.0-rc2. Packages with Linux kernel 6.1, 6.6 and 6.12 are available for installation.

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://forum.rosa.ru/viewtopic.php?t%3D11359

Week 10

Ubuntu Touch OTA-8 Focal Released:

03/03/2025

The OTA-8 Focal (over-the-air) firmware has been released, developed by the UBports project, which took over the development of the Ubuntu Touch mobile platform after Canonical stepped away from it. This is the eighth release of Ubuntu Touch, based on the Ubuntu 20.04 package base. The project also develops an experimental port of the Unity 8 desktop, which has been renamed Lomiri.

Ubuntu Touch OTA-8 Focal update will be generated in the coming days for Asus Zenfone Max Pro M1, F(x)tec Pro1 X, Fairphone 3/3+/4, Google Pixel 3a/3a XL, JingPad A1, Oneplus 5/5T/6/6T, OnePlus Nord N10 5G/N100, Sony Xperia X, Vollaphone X/22/X23 and Xiaomi Poco X3 NFC / X3. Compared to the previous release, builds have been added for Xiaomi Poco M2 Pro, Xiaomi Redmi Note 9 Pro/Pro Max/9S, Volla Phone Quintus and Volla Tablet.

The development of the new release focused on preparing for the transition to the newer LTS branch of Ubuntu. One of the big changes in Ubuntu Touch OTA-8 Focal:

Support for VoLTE (Voice over LTE), a technology for transmitting voice over LTE networks, has been implemented. Using VoLTE is currently only possible on Volla Phone X23 and Volla Phone 22 smartphones, the firmware for which uses Halium 12 system components.

https://ubports.com/en/blog/ubports-news-1/post/ubuntu-touch-ota-8-focal-release-3953

OpenBot 0.8 Creation Platform Released:

04/03/2025

After almost two years of development, the OpenBot 0.8 open platform has been released. It is designed to create moving wheeled robots based on a regular Android smartphone. The platform was created in the research division of Intel and develops the idea of using the computing capabilities of a smartphone and the built-in GPS, gyroscope, compass and camera in the creation of robots. The code, written in Swift and Java, is distributed under the MIT license.

The software for robot control, environment analysis and autonomous navigation is implemented as an application for the Android platform. It is assumed that the platform can be useful for teaching robotics, quickly creating your own prototypes of moving robots and conducting research related to autopilots and autonomous navigation.

OpenBot allows you to start experimenting with moving robots at a minimum cost - to create a robot, you can get by with a mid-range smartphone and additional components totaling about USD50. The chassis for the robot, as well as related parts for attaching the smartphone, are printed on a 3D printer according to the proposed layouts. If you do not have a 3D printer, you can cut a frame out of cardboard or plywood. Movement is provided by four electric motors.

To control the engines, attachments and additional sensors, as well as to monitor the battery charge, an Arduino Nano board, based on the ATmega328P microcontroller, is used, which is connected to the smartphone via a USB port. Additionally, connection of speed sensors and ultrasonic sonar is supported. Remote control of the robot can be carried out via a client application for Android, via a computer located on the same WiFi network, via a web browser or via a game controller with Bluetooth support (for example, PS4, XBox and X3).

https://github.com/ob-f/OpenBot/releases/tag/v0.8.0

Release of Godot 4.4 game engine:

04/03/2025

After seven months of development, the free Godot 4.4 game engine, suitable for creating 2D and 3D games, has been released. The engine supports an easy-to-learn game logic language, a graphical environment for designing games, a one-click game deployment system, animation and physics simulation capabilities, a built-in debugger, and a performance bottleneck detection system. The code for the game engine, game design environment, and related development tools (physics engine, sound server, 2D/3D rendering backends, etc.) is distributed under the MIT license.

As a result of a conflict in the community, where Godot had doubled down on retarded gender politics, a fork of the project was founded last year - Redot. The fork was a response to the strict moderation policy in official communication channels and the rejection by Godot developers of opinions that differ from their views. The reasons for creating the fork were the desire to solve problems with community management and return to work for the benefit of the community, and not for corporate interests. The Redot project is trying to build an inclusive community in which politicization of the development process is unacceptable, and the main focus is only on creating games. Currently, the fork's development activity lags slightly behind the Godot engine.

https://godotengine.org/article/godot-4-4-a-unified-experience/

Thunderbird 136.0 Mail Client Released:

05/03/2025

Thunderbird 136.0, a community-developed email client based on Mozilla technologies, has been released. Thunderbird 136 is built on the Firefox 136 codebase and is categorized as an intermediate version, with updates released before the next release. Thunderbird 128.8.0 has been released in the ESR branch, which has a long-term support period and updates are released throughout the year.

The publication of Thunderbird 136.0 marked a change in the approach to forming releases. In addition to the previous

scheme, in which a major release was formed once a year and based on the ESR branches of Firefox, major releases will now also be published once a month and based on the Firefox releases with a regular maintenance period. Previously, builds based on regular Firefox releases were not reflected on the download page, positioned as testing, and provided with a special warning in the release notes. Thunderbird 136 is offered by default and without the label of being intended for testing only. The formation and maintenance of the Thunderbird ESR branches will continue without changes.

It is expected that the transition to monthly major releases will speed up the delivery of new features to users, smooth the transition to new major releases (changes will be spread out over time), and expand the range of bugs fixed with each update (the ESR branch only fixed vulnerabilities and serious issues).

https://blog.thunderbird.net/2025/03/thunderbird-release-channel-update/

Intel Adds eUSB2V2 Support to Linux Kernel:

05/03/2025

An engineer from Intel has prepared a change to support the second version of the eUSB2 extension (eUSB2V2 - Embedded USB2 Version 2.0). The change has been transferred to the usb-next branch, which tests the Linux kernel functionality planned for release 6.15 related to USB.

eUSB2V2 is an extension of the USB 2.0 standard that reduces the supply voltage (to 1.2 volts) and increases performance. The data transfer rate in eUSB2V2 can reach 4.8 Gbps, which is 10 times faster than the usual 480 Mbps typical for USB 2.0. The proposed technology allows laptop manufacturers to equip their devices with higher-resolution webcams, while continuing to use the Embedded USB2 bus for connecting them.

https://web.git.kernel.org/pub/scm/linux/kernel/git/gregkh/usb.git/commit/?h%3Dusb-next%26id

%3Dc749f058b4371430a8338e1ca72b9ae38fef613b

FerretDB 2.0 Release:

05/03/2025

FerretDB 2.0 is now available, allowing you to replace the proprietary document-oriented MongoDB DBMS with a fully open source software stack based on PostgreSQL without changing your application code. The code is written in Go and is licensed under Apache 2.0. The release is marked as ready for production deployments.

FerretDB's primary target audience is MongoDB users who want to use a fully open source software stack. FerretDB supports a subset of MongoDB features that are most commonly used in typical applications. The need to implement FerretDB may arise due to MongoDB's transition to the proprietary SSPL license, which is based on the AGPLv3 license but is not open source because it contains a discriminatory requirement to distribute under the SSPL license not only the application code itself, but also the source code of all components involved in providing the cloud service.

The key change in FerretDB 2.0 is the transition to DocumentDB, a Microsoft-opened extension to PostgreSQL that implements the ability to store data in a JSON-like BSON (Binary JSON) format compatible with MongoDB. Initially, FerretDB's work was limited to translating MongoDB requests into SQL queries to PostgreSQL. The transition to DocumentDB has increased FerretDB's performance by more than 20 times for some types of workload. Other new features in FerretDB 2.0 include improved compatibility with MongoDB, replication , and support for vector search .

https://blog.ferretdb.io/ferretdb-v2-ga-open-source-mongodb-alternative-ready-for-production/

Release of FEX 2503 emulator:

06/03/2025

The FEX 2503 project has been published. It develops an emulator for running games and applications built for x86 and x86-64 architectures in a Linux environment on systems with ARM64 (AArch64) processors. The FEX emulator is used by the Asahi project to run games from the Steam catalog built for the x86_64 architecture on systems with ARM chips. The project code is written in C++ with assembler inserts and is distributed under the MIT license.

Libraries required to run x86 applications in the ARM64 environment are connected in the form of an overlay layer with the root file system image (rootfs), supplied in the SquashFS format. Such an overlay makes it possible to do without creating a separate chroot environment. Loading of rootfs images already generated by the project is performed by the FEXRootFSFetcher utility.

To access host environment capabilities, such as audio and 3D graphics tools, thunk libraries are placed in the rootfs, which translate calls to libraries and code on the host system side (for example, calls to OpenGL and Vulkan can be redirected). Such libraries can also be used on the host side, to call code in the guest environment. Among the libraries available for forwarding are: libEGL, libGL, libSDL2, libX11, libVDSO, libasound, libdrm, libfex_malloc, libvulkan, libwayland-client and libxshmfence.

https://fex-emu.com/FEX-2503/

Google Publishes Toolkit for Analyzing and Modifying AMD Microcode:

06/03/2025

Engineers from Google have disclosed details of a vulnerability (CVE-2024-56161) that allows bypassing the digital signature verification mechanism when updating microcode in AMD processors based on the 1st to 4th generations of the Zen microarchitecture. At the same time, the Zentool toolkit, developed during the study of methods for working with microcode in AMD processors, was published under the Apache

2.0 license. A guide to the RISC86 microarchitecture used in AMD microcode and a note on creating your own microcode were also prepared. It shows how you can create your own processor instructions implemented on RISC86 microcode, change the behavior of existing instructions, and load microcode changes into the processor.

For simplification, the Zentool toolkit has been developed, which allows you to analyze microcode, manipulate microcode, and create patches that can be used to change the microcode in AMD Zen processors. In the future, the developers plan to continue expanding the capabilities of the toolkit and provide a kind of analogue of binutils, but not for traditional machine code, but for microcode.

Loading custom patches for AMD Zen 1-4 processor microcode was made possible by a vulnerability caused by the use of the CMAC impersonation algorithm during verification instead of a strong hash function. AMD fixed the vulnerability in a December microcode update by replacing CMAC with a cryptographically strong hash function.

To digitally sign the microcode loaded into the processor, AMD uses a private RSA key, and adds a public key as part of the microcode patch. To verify that the public key matches the original RSA key pair, the processor compares the hash of the AMD public key embedded in the CPU during manufacturing with the hash of the public key specified in the patch.

AMD uses a single encryption key for AES-CMAC, which is shipped on all CPUs from Zen 1 to Zen 4. Thus, it is enough to extract this key from any AMD CPU and it will be applicable to all other CPUs. The researchers found that the Zen 1 through Zen 4 generations of processors use a known key for AES-CMAC encryption, taken from an example mentioned in the NIST SP 800-38B block cipher recommendation .

https://bughunters.google.com/blog/5424842357473280/zen-and-the-art-of-microcode-hacking https://github.com/google/security-research/blob/master/pocs/cpus/entrysign/zentool/docs/reference.md

Blender-Made Cartoon 'Flow' Wins Oscar:

06/03/2025

The 2025 Academy Award for Best Animated Feature Film has been awarded to the animated film Flow, directed by Latvian Gints Zilbalodis. The film follows the adventures of a cat during a flood and is notable for being created by an independent team on a limited budget in the free 3D modeling system Blender. Work on the film lasted five and a half years.

In his previous works, the director used the Maya 3D package, but in 2019 he switched to Blender, after the EEVEE rendering engine was released in Blender 2.80. EEVEE supported real-time rendering and allowed him to implement the desired workflow. One of the most important selection criteria was speed, not only of rendering, but also of working with files, setting up lighting, and creating a general style. Blender also had all the tools the director needed.

https://en.wikipedia.org/wiki/Flow_(2024_film)

Samba 4.22.0 Released:

06/03/2025

After 6 months of development, Samba 4.22.0 was released, continuing the development of the Samba 4 branch with a full implementation of a domain controller and Active Directory service, compatible with the Windows 2008 implementation and capable of servicing all Microsoft-supported versions of Windows clients, including Windows 11. Samba 4 is a multifunctional server product, also providing the implementation of a file server, print service and identification server (winbind).

http://www.samba.org/

PipeWire 1.4.0 Released:

07/03/2025

After eight months of development, a new stable branch of the PipeWire 1.4.0 multimedia server is out. It has replaced the PulseAudio sound server and differs from it by adding tools for working with video streams, the ability to process sound with minimal delays and a new security model for managing access at the level of individual devices and streams. The project is supported in GNOME and is used by default in Fedora Linux, RHEL, Ubuntu, Debian, SUSE/openSUSE and many other Linux distributions. The project code is written in C and distributed under the MIT license.

PipeWire is based on a multi-process architecture that enables content sharing between multiple applications. It provides capabilities for handling any media streams, mixing and redirecting video streams, and managing video sources such as video capture devices, webcams, or screen content output by applications. PipeWire enables multiple applications to work together with a webcam and solves the problems of secure screen capture and remote screen access in a Wayland environment.

When used as an audio server, PipeWire can provide low latency and functionality that combines the capabilities of PulseAudio and JACK, including the needs of professional audio systems that PulseAudio could not claim. PipeWire offers an advanced security model that allows access control at the level of individual devices and specific streams. The implemented access model simplifies the routing of audio and video in and out of isolated containers.

https://gitlab.freedesktop.org/pipewire/pipewire/-/releases/1.4.0

Apertis 2025.0:

08/03/2025

Collabora has introduced the Apertis 2025.0 Linux distribution, originally created for automotive systems, but then reoriented for a wider range of electronic devices, embedded equipment and industrial equipment. Examples of devices where Apertis is used include the Atari VCS gaming console, Raspberry Pi 4

boards, R-car automotive SoCs and the Bosch D-tect 200 wall scanner .

Reference system images are distributed for the x86_64, arm64 and armhf architectures. The distribution is modular and allows device manufacturers to independently form the necessary system environment filling. Formation of builds based on traditional deb packages and monolithic atomically updated images based on OSTree are supported. The maintenance period for each Apertis release is 1 year and 9 months, every three months a corrective release with error corrections is built. The distribution is based on packages from Debian GNU/Linux 12.

Apertis allows you to create builds that do not include code under the GPLv3 family of licenses. Instead of using outdated versions of GNU utilities created before the transition to the GPLv3 license, Apertis uses more modern alternatives under permissive licenses. For example, instead of the GNU coreutils and findutils packages, Apertis offers analogs from the uutils project, written in the Rust language and distributed under the MIT license, and instead of GnuPG, Sequoia-PGP is supplied under the GPL-2+ and LGPL-2+ licenses. For those who do not care about the legal issues associated with GPLv3, the option of using traditional utility sets is left.

The Apertis project adheres to the Debian development rules and includes only applications that are supplied under open licenses or allow free distribution. In order for companies creating their products based on Apertis to be sure of the license purity of derivative works, a SBOM (Software Bill of Materials) report is generated for each build, which indicates information about the licenses of all used code files, as well as data on program versions, which is also convenient for checking for vulnerable versions.

https://www.collabora.com/news-and-blog/news-and-events/apertis-v2025-the-second-bookworm-based-release-for-industrial-embedded-devices.html

Linux terminal for Google Pixel devices:

08/03/2025

Google included the Linux Terminal app in the March Android 15 firmware update ("Pixel Feature Drop") for Google Pixel devices. The app allows you to run a Debian GNU/Linux virtual machine in the Android environment, where you can run regular Linux applications.

Linux Terminal is activated in the developer settings (section "Settings > System > Developer > Linux development environment"). For this option to appear, you need to activate the developer mode by quickly pressing the build number seven times on the "Settings → About Phone" page. After the first launch, the application offers to download a virtual machine image with Linux to the device, which takes up about 500 MB.

The Terminal application is being developed in the AOSP (Android Open Source Project) repository in the core Android platform. The functionality of the Linux virtual machine is being developed within the Ferrochrome project. Debian GNU/Linux 12 is running in the guest environment. The AVF (Android Virtualization Framework) framework is used for virtualization, implemented, based off the KVM hypervisor and the crosvm toolkit .

https://www.androidpolice.com/android-15-linux-terminal-app/

Week 11

Nova driver starter code written in Rust is available for Linux kernel 6.15:

10/03/2025

A set of patches with an initial implementation of the Nova driver for NVIDIA GPUs written in Rust has been proposed for inclusion in the Linux 6.15 kernel, which is expected to be released in late May. At the first stage, the nova-core framework, which contains about 400 lines of code and implements a basic level of abstractions over the GSP firmware software interfaces, was transferred to the kernel. In addition to nova-core, the patch includes some bindings necessary for the drivers to work with the firmware.

At the next stage, the kernel plans to include the nova-drm (Direct Rendering Manager) DRM driver for interaction with the GPU from user space, as well as the VFIO driver with the vGPU manager , which allows the use of NVIDIA virtual GPUs in virtualization systems.

The Nova driver is designed for use with NVIDIA GPUs equipped with GSP firmware, which is used starting with the NVIDIA GeForce RTX 2000 series based on the Turing microarchitecture. In such GPUs, initialization and control operations are implemented in the firmware and are performed by a separate GSP (GPU System Processor) microcontroller. Nova is expected to replace the Nouveau driver for GPUs with GSP support in the long term.

The project is being developed by Red Hat employees with the goal of getting rid of the complications inherent in the Nouveau DRM driver. Significant simplification is achieved by using ready-made handlers provided by the GSP firmware and dropping support for older GPUs. Nova will also solve architectural problems that require significant reworking of the code base in Nouveau, such as problems with locks in the

VMM/MMU code.

In addition to Nova, drivers for Apple AGX GPU (drm-asahi), NVMe (rnvme) and Android Binder are being developed in Rust for the Linux kernel. Cisco is developing the PuzzleFS file system for the Linux kernel in Rust . The QR code generator for the emergency stop screen, the rnull block device driver (replacement for null_blk) and the ax88796b_rust driver for the PHY interface of the ASIX AX887xx Ethernet controller have already been accepted into the kernel.

https://lore.kernel.org/all/Z84dHHEn6xfvlRxk@cassiopeiae/T/%23u

FreeBSD 13.5 ReleaseD:

11/03/2025

After 6 months of development, FreeBSD 13.5 release has been published, which will be the last in the FreeBSD 13.x branch. Updates for FreeBSD 13.5 will be released until April 30, 2026. In parallel, the FreeBSD 14 branch is being developed, the next release (14.3) of which is scheduled for June 3, 2025. The first release of the FreeBSD 15 branch will be formed in December 2025.

FreeBSD 13.5 installation images are generated for the amd64, i386, powerpc, powerpc64, powerpc64le, powerpcspe, armv6, armv7, aarch64 and riscv64 architectures. Additionally, builds have been prepared for virtualization systems (QCOW2, VHD, VMDK, raw) and cloud environments Amazon EC2, Google Compute Engine and Vagrant.

https://www.freebsd.org/releases/13.5R/announce/

Haiku Ports Latest Firefox, LibreWolf, and Thunderbird:

11/03/2025

Gerasim Troeglazov (3dEyes) has compiled fresh versions of the Firefox 136 browser, as well as the LibreWolf 136 browser

based on this branch and the ThunderBird 136 email client.

According to one Haiku developer, BeOS enthusiasts were among the inspirations for Firefox in the early 2000s. At the time, the Bezilla project was developing a port of the Mozilla Suite for BeOS. Since the suite was too bloated, the BeOS community tried to create a lightweight version based on it, which would leave only the browser and remove all other components, such as the email client and web page editor. Mozilla developers found the idea worthy of attention and released their own standalone version of the browser called Phoenix, later renamed Firebird due to overlap with the trademark, and then renamed Firefox again due to overlap with the name of a free DBMS.

https://t.me/haiku_ru/151488

Mesa replaced the Nouveau OpenGL driver with Zink for new NVIDIA GPUs:

11/03/2025

Collabora has announced that a change has been adopted into the Mesa codebase that replaces the default OpenGL driver for NVIDIA GPUs starting with the Turing microarchitecture. The next release of Mesa 25.1 will replace the Nouveau (nvc0) OpenGL driver for GPUs with the Collabora-developed Zink OpenGL driver in conjunction with the NVK Vulkan driver. Compared to Nouveau, the Zink driver demonstrates higher performance in many tests and is not susceptible to the problems that Nouveau has when running on new NVIDIA GPUs.

Zink provides an implementation of OpenGL 4.6 on top of Vulkan, enabling hardware-accelerated OpenGL on devices that support the Vulkan API. Zink's performance is close to that of native OpenGL implementations, allowing you to focus on providing high-quality Vulkan API support and implement OpenGL support on top of Vulkan instead of spending resources on maintaining separate OpenGL drivers. NVIDIA's Vulkan implementation is based on the NVK driver, which supports

Vulkan 1.4 for NVIDIA Turing (GeForce GTX 16xx, RTX 20xx, and Quadro RTX series), Ampere (GeForce RTX 30xx and RTX A2000/4000/5000/6000 series), Ada (GeForce RTX 4xxx, RTX 4000 SFF, RTX 4xxx/5000/6000 Ada series), and newer microarchitectures.

https://www.collabora.com/news-and-blog/news-and-events/goodbye-nouveau-gl-hello-zink.html

CrossOver 25.0 Released:

12/03/2025

After a year of development, CodeWeavers has released the Crossover 25.0 package, based on Wine code and designed to run programs and games written for the Windows platform. CodeWeavers is one of the key participants in the Wine project, sponsors its development and returns to the project all the innovations implemented for its commercial products. The source code of the open components of CrossOver 25.0 can be downloaded from their web page .

https://www.codeweavers.com/support/forums/announce/?t%3D24;msg%3D322440%0D%0A

GStreamer 1.26.0 Available:

12/03/2025

After a year of development, GStreamer 1.26 is released, a cross-platform set of components for creating a wide range of multimedia applications, from media players and audio/video file converters to VoIP applications and streaming systems. The GStreamer code is distributed under the LGPLv2.1 license. Updates are also being developed for the gst-plugins-base, gst-plugins-good, gst-plugins-bad, gst-plugins-ugly plugins, as well as the gst-libav binding and the gst-rtsp-server streaming server. At the API and ABI level, the new release is backward compatible with the 1.0 branch. Binary builds will soon be prepared for Android, iOS, macOS and Windows (on Linux, they recommend you use packages from the distribution).

https://lists.freedesktop.org/archives/gstreamer-devel/2025-March/082142.html

The United Nations has proposed Open Source principles:

12/03/2025

The Open Source Initiative (OSI), an organization that reviews licenses for open source compliance, has formally endorsed the Open Source Principles, proposed by the United Nations (UN) to encourage collaboration and widespread adoption of open source code within UN entities and elsewhere. The principles were developed by Open Source United, a community organized by the UN Executive Board's Digital Network program, and provide a framework for the use, development, and dissemination of open source code throughout the organization.

https://opensource.org/blog/osi-endorses-united-nations-open-source-principles

Ubuntu 25.10 to replace GNU Coreutils with uutils written in Rust:

13/03/2025

Jon Seager, Canonical's vice president of engineering and the technical lead for the Ubuntu project, has unveiled an initiative to replace Ubuntu's system utilities with Rust-based ones. The initiative's first goal is to move Ubuntu 25.10 to using the uutils toolkit by default instead of the GNU Coreutils suite. If the experiment is successful, uutils will also be used by default in the Ubuntu 26.04 LTS branch.

The replacement will affect more than a hundred utilities included in Coreutils, including sort, cat, chmod, chown, chroot, cp, date, dd, echo, hostname, id, ln and ls. Currently, uutils utilities are already used by default in the Debian-based Apertis distribution , as well as in the independent AerynOS (SerpentOS) distribution. Last week's release of the uutils

coreutils package successfully passes 507 tests (506 in the previous release, 476 in the previous one) from the GNU Coreutils benchmark test suite. 69 tests failed, and 41 tests were skipped. In the coming weeks, work is also planned to begin on replacing the su and sudo utilities in Ubuntu with the sudo-rs package. Of the projects under consideration, zlib-rs and ntpd-rs are also mentioned.

The reason for the migration is said to be the desire to improve the reliability and security of the utilities that underlie the distribution. Using Rust will reduce the risk of errors when working with memory, such as accessing a memory area after it has been freed and going beyond the buffer boundaries. According to John Seeger, protection against these errors will increase security guarantees, and with increased security, the overall reliability of the system will increase.

It is noted that Canonical is considering various methods of improving quality, and one of them is the delivery of programs that are initially developed with an eye on security, reliability, and correctness. This is especially important for the basic components of the distribution, since if problems arise in low-level software, these problems are reflected in the work of all higher layers, for example, if there are problems with performance in the basic packages, they affect the performance of other subsystems.

To test replacing system components in Ubuntu, the oxidizr project has been prepared, offering a command-line toolkit for managing system experiments related to replacing traditional utilities with alternatives written in Rust. Currently, oxidizr offers experiments for switching to the default use of the uutils coreutils, uutils findutils, uutils diffutils and sudo-rs packages. For example, to replace coreutils and findutils in your system, it is enough to run the command "sudo oxidizr enable --experiments coreutils findutils", and to return to the original state, you can use the command "oxidizr disable".

https://discourse.ubuntu.com/t/carefully-but-purposefully-oxidising-ubuntu/56995

Cozystack Project Accepted into CNCF:

14/03/2025

The technical committee of the CNCF (Cloud Native Computing Foundation), part of the Linux Foundation, approved the adoption of the Cozystack project, which develops a platform for building private clouds and PaaS. Cozystack has received the status of a CNCF sandbox project. The developers of Cozystack have begun the process of integration with the CNCF infrastructure.

The Cozystack platform is distributed under the Apache 2.0 license and allows you to build a cloud on your existing hardware for deploying "cloud native" and open source tools: managed Kubernetes clusters, DBaaS (database as a service), SaaS (application as a service) and virtual machines based on KubeVirt. Cozystack also provides a ready-made stack for monitoring and alerting built using Victoria Metrics, Victoria Logs, Grafana and Alerta. Cozystack can be used to provide services managed by Kubernetes, build geo-distributed clusters and organize the operation of databases deployed on separate equipment (bare metal).

The non-profit organization CNCF (Cloud Native Computing Foundation) is part of the Linux Foundation and oversees cloud native projects such as Kubernetes, Etcd, Envoy, Prometheus, Cilium, Istio, K3s and FluxCD. To obtain the status of a CNCF main project, a project must go through the "Sandbox", "Incubating" and "Graduated" stages. CNCF Sandbox is a kind of entry point for projects that want to join CNCF and become part of it.

The project's move to CNCF gives Cozystack users a guarantee that the platform will always be available under the Apache 2.0 license and will not suffer the fate of projects such as Mongo, Redis, Terraform and Vault, whose licenses were changed to proprietary. In addition, the project's inclusion in CNCF, it will attract new developers and users, and make the project's management more transparent.

https://github.com/cncf/sandbox/issues/322%23issuecomment-2697791780

Release of WebKitGTK 2.48.0:

14/03/2025

The release of a new stable branch WebKitGTK 2.48.0 is presented, a port of the WebKit browser engine for the GTK platform. WebKitGTK allows you to use all the capabilities of WebKit through a GNOME-oriented programming interface based on GObject and can be used to integrate web content processing tools into any applications, from using in specialized HTML/CSS parsers to creating full-featured web browsers. Among the well-known projects using WebKitGTK, there is the standard GNOME browser (Epiphany). Previously, WebKitGTK was used in the Midori browser, but after the project was transferred to the Astian Foundation, the old version of Midori on WebKitGTK was abandoned and a fundamentally different product was created with the same name Midori.

https://webkitgtk.org/2025/03/14/webkitgtk2.48.0-released.html

GTK 4.18 is available:

15/03/2025

After six months of development, the multi-platform GUI toolkit, GTK 4.18 has been released. GTK 4 follows a development process that tries to provide application developers with a stable and supported API for several years that they can use without having to rework their applications every six months because of API changes in the next GTK branch.

Among the most notable improvements in GTK 4.18 was an experimental backend has been added that allows running GTK applications on Android smartphones. The ability to use OpenGL for rendering on the Android platform has been implemented. The backend for the X11 protocol has been deprecated and is set to be discontinued in the GTK 5 branch.

The deprecation is explained by the cessation of activity on the X11 protocol development and problems with maintainers - the backend is supported on a residual basis, since the existing GTK and GNOME developers are focused on Wayland. Due to stagnation in the backend development, it slows down the implementation of new features in GTK.

https://gitlab.gnome.org/GNOME/gtk/-/tags/4.18.0

Debian 12.10 Release. First Stage of Debian 13 Freeze:

16/03/2025

The tenth corrective update of the Debian 12 distribution has been generated , which includes accumulated package updates and fixes to the installer. The release includes 66 updates with fixes for stability issues and 43 updates with fixes for vulnerabilities.

Among the changes in Debian 12.10, we can note the update to the latest stable versions of the bup, intel-microcode, mariadb, postgresql-15, spamassassin, systemd and tzdata packages. The kanboard and libnet-easytcp-perl packages, which remained unmaintained and had security issues, were removed, as well as the looking-glass package, which was deemed unsuitable for a stable release.

Additionally, it is worth noting that the Debian 13 "Trixie" branch has been moved to the first stage of freezing the package base. At the first stage of freezing, "transitions" (updating packages that require adjusting dependencies in other packages, which leads to temporary removal of packages from Testing) have been stopped, and updating packages necessary for building (build-essential) has been stopped.

On April 15, 2025, a soft freeze of the package base will occur, stopping the acceptance of new upstream packages and preventing the re-inclusion of previously removed packages. On May 15, 2025, a hard freeze will be applied before the release, stopping the process of moving core packages and packages without autopkgtests from the unstable branch to testing

completely, and beginning a phase of intensive testing and fixing of release-blocking issues. Some time after the hard freeze, a full freeze will be applied, covering all packages. Debian 13 is expected to be released in the second half of 2025.

https://www.debian.org/News/2025/20250315

Week 12

Release of GIMP 3.0.0:

17/03/2025

Seven years after the release of the 2.10 branch, the GIMP 3.0 graphics editor has been released. Ready-made builds have been published for Linux (AppImage and Flatpak for x86_64 and ARM64). Builds for Windows and macOS are being prepared for publication.

In preparation for future releases, it has been decided to move to a more predictable and frequent schedule of publishing new stable branches. In the future, the developers will not try to push many major changes at once, but will try to focus on polishing individual new features. The next major branch of GIMP 3.2, is planned to be published in about a year. In between major releases, corrective updates with bug fixes will be released.

Among the most significant improvements in GIMP 3.0 :

The transition to the GTK3 library, CSS-like style definition system and the use of client-side dialog window decoration (CSD, window title and frames are drawn not by the window manager, but by the application itself) was made. New widgets were proposed. Also, providing native support for working in environments based on the Wayland protocol. Then they added support for HiDPI and taking into account system scaling settings when generating the interface. The interface has been significantly modernized. Support for symbolic icon sets has been improved, and now automatically adjusts to the set foreground and background colors (when switching from light to dark mode, you no longer need to manually change the icon set). And many more!

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://github.com/GNOME/gimp/

releases/tag/GIMP_3_0_0

Introducing TinyKVM for process-level virtualization:

17/03/2025

Varnish Software, a company that develops systems for building content delivery and caching networks, has introduced the open source TinyKVM project, that develops tools for isolating the execution of individual processes using the KVM hypervisor. The stated goal of the project is to create the fastest sandbox-isolation system for individual processes using hardware virtualization. The project code is written in C and C + + and is distributed under the GPLv3 license (for those who are not ready to comply with the requirements of GPLv3, a commercial license is provided).

TinyKVM is designed to launch, isolated, any console programs for Linux with performance close to normal execution. The overhead for calling each system call is about 2 microseconds. Additional isolation of processes in caching systems and web request processing is mentioned as an example of the project's application. TinyKVM is designed to replace the libriscv emulator used to isolate the processing of each web request in the Varnish platform. Additionally, a version of the libvmod library has been created, allowing the execution of Varnish modules using TinyKVM.

When running with TinyKVM, the machine code of programs is executed without emulation layers on the CPU and is limited using the KVM hypervisor API, which eliminates overhead and achieves performance close to that of running without virtualization

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://info.varnish-software.com/blog/tinykvm-the-fastest-sandbox

New stable version of Vivaldi 7.2:

19/03/2025

The proprietary browser Vivaldi 7.2 has been released. It is being developed on the Chromium engine by former developers of the Opera Presto engine. Vivaldi builds are available for Linux, Windows, and macOS. The project distributes changes made to the Chromium code base under an open license. The browser interface is written in JavaScript using the React library, Node.js platform, Browserify, and various ready-made NPM modules. The interface implementation is available in source code, but under a proprietary license.

The project aims to create a customizable and functional browser that preserves the privacy of user data. The main functions include a tracking and advertising blocker, note, history and bookmark managers, private browsing mode, end-to-end encrypted synchronization, tab grouping mode, sidebar, configurator with a large number of settings, horizontal tab display mode, and a built-in mail client, RSS reader and calendar in test mode.

https://vivaldi.com/blog/vivaldi-on-desktop-7-2/

Release of elementary OS 8.0.1:

19/03/2025

The release of elementary OS 8.0.1 has been published. It is positioned as a fast, open, and privacy-conscious alternative to Windows and macOS. The project focuses on high-quality design aimed at creating an easy-to-use system that consumes minimal resources and provides high startup speed. Users are offered their own Pantheon desktop environment. Bootable iso images (3.3 GB) are available for download, available for the amd64 architecture (for a free download from the project's website, enter 0 in the donation amount field).

When developing the original components of elementary OS, GTK3, the Vala language and its own Granite framework are used. The developments of the Ubuntu project are used as the base of the distribution. At the packages and repository support level, elementary OS 8.x is compatible with Ubuntu 24.04. All additional applications offered for installation through

AppCenter, as well as some applications supplied by default, are packaged using the Flatpak format. The graphical environment is based on its own Pantheon shell, which combines such components as the Gala window manager (based on LibMutter), the Slingshot launcher, the Switchboard control panel, the Wing top panel, the Dock taskbar and the Pantheon Greeter session manager (based on LightDM).

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://blog.elementary.io/os-8-0-1-available-now/

SDL3 library with support for older systems:

19/03/2025

The author of the LDL (Little Directmedia Layer) graphics library and the ArcanumWorld game engine is developing the SDL3Lite project, aimed at recreating the SDL3 library with support for older architectures and operating systems, while maintaining compatibility with SDL3 at the C API level. With limited functionality, SDL3Lite already supports Linux, starting with Debian 3, and Windows 95. The system requirements for the graphics API include support for OpenGL 1.0, and for 2D rendering - OpenGL 1.2. In the future, the author plans to add support for Windows 3.1 and MS-DOS. The code is written in C + + and is distributed under the BSL (Boost Software License).

https://github.com/JordanCpp/SDL3Lite

Ubuntu rebuilding performance increased:

19/03/2025

The results of the performance impact assessment of rebuilding packages for Ubuntu with different options and implementations of memory allocation functions have been published. The experimenter managed to increase the performance of the jq package with tools for processing data in the JSON format by 90% (1.9 times) by simply rebuilding it from the same package with the source code, without making

any changes to the code itself. The performance was assessed by measuring the execution time of a typical filtering query over GeoJSON data, 500 MB in size.

Results of the experiment:

The version compiled in GCC from the same source code with default flags turned out to be 2-4% faster than the Ubuntu binary package. Rebuilding in Clang 18 with optimization level "-O3", enabling link-time optimizations ("-flto") and disabling debug information ("-DNDEBUG") resulted in a 20% speedup.

Rebuilding with the TCMalloc memory allocation system (adding "-L/usr/lib/x86_64-linux-gnu -ltcmalloc_minimal" to LDFLAGS) resulted in a 40% speedup.

Replacing malloc functions with tcmalloc , jemalloc and mimalloc memory allocation systems via "LD_PRELOAD = /usr/lib/x86_64-linux-gnu/lib....so" resulted in 27%, 29% and 44% performance gains. When running with mimalloc, which showed a 44% speedup, the environment variable "MIMALLOC_LARGE_OS_PAGES = 1" was set.

Rebuilding the package with mimalloc in LDFLAGS instead of linking via LD_PRELOAD resulted in a 90% speedup in the test. Another test of processing 2.2GB of JSON data in 13,000 files also showed a roughly twofold performance increase.

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://gist.github.com/jwbee/7e8b27e298de8bbbf8abfa4c232db097

GNOME DE Release 48:

19/03/2025

After six months of development, the GNOME 48 desktop environment has been released. Specialized Live builds based on openSUSE and an installation image prepared by the GNOME OS initiative are available for a quick evaluation of GNOME 48's capabilities. GNOME 48 is also already included

in experimental builds of Ubuntu 25.04 and Fedora 42.

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://foundation.gnome.org/2025/03/19/introducing-gnome-48/

Release of miracle-wm 0.5:

20/03/2025

Matthew Kosarek, a developer from Canonical, has released the miracle-wm 0.5 compositing manager, which uses the Wayland protocol and Mir compositing manager components. Miracle-wm supports tiling window layout, similar to the i3 and Sway projects. Waybar can be used as a panel. The project code is written in C++ and is distributed under the GPLv3 license. Ready-made builds are available in the snap format, as well as in rpm and deb packages for Fedora and Ubuntu .

The goal of miracle-wm is to create a composite server that uses tiled window management, but is more functional and stylish than products like Swayfx. At the same time, the project allows you to use classic techniques for working with floating windows, such as placing individual windows on top of a tiled grid or pinning windows to a specific place on the desktop. Virtual desktops are supported with the ability to set a default window mode for each desktop (tiled layout or floating windows).

It is expected that miracle-wm will be useful for users who prefer a tiled layout, but want visual effects and a more vivid graphical design with smooth transitions and colors. The configuration is defined in YAML format. To install miracle-wm, you can use the command "sudo snap install miracle-wm --classic".

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://github.com/mattkae/miracle-wm/releases/tag/v0.5.0

Release of LosslessCut 3.65.0:

20/03/2025

LosslessCut 3.65.0 has been released. It provides a graphical interface for editing multimedia files without recoding the content. The most popular function of LosslessCut is cropping and trimming video and audio, for example, to reduce the size of large files shot on an action camera or a quadcopter camera. LosslessCut allows you to select relevant fragments of the recording in the file and discard the excess, without performing a full recoding and preserving the original quality of the material. Since the processing is performed by copying the existing data, and not recoding, the operations are very fast. LosslessCut is written in JavaScript using the Electron platform and is an add-on to the FFmpeg package. The code is distributed under the GPLv2 license. Builds are prepared for Linux (snap , flatpak), macOS and Windows.

Without recoding, the program can also solve tasks such as attaching an audio track or subtitles to a video, cutting out individual scenes from videos (for example, cutting out commercials from TV program recordings), saving fragments linked to tags/chapters separately, regrouping video parts, dividing audio and video into different files, changing the type of multimedia container (for example, from MKV to MOV), saving individual video frames as images, creating thumbnails, exporting a fragment to a separate file, changing metadata (for example, location data, recording time, horizontal or vertical orientation). There are tools for detecting and automatically cutting out empty areas (black screen in video and fragments without sound in audio files), as well as linking to scene changes.

You can combine fragments from different files, but the files must be encoded using the same codec and parameters. It is possible to edit individual parts with selective recoding of only the changed data, but leaving the rest of the information in the original video that was not affected by editing. During editing, rollback of changes (undo/redo) and display of the FFmpeg command log are supported (you can repeat typical operations from the command line without using LosslessCut).

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://github.com/mifi/lossless-cut/releases/tag/v3.65.0

NVIDIA driver and GIMP 3.0 for Haiku OS:

21/03/2025

Ilya Chugin (X512) reported on the first results of porting the NVIDIA driver for the Haiku operating system. The port is still in its early stages of development, but is already suitable for simple rendering. The work involves an open kernel module from NVIDIA and the Vulkan driver NVK from Mesa. The Mesa driver NVK has been supplemented with support for the kernel-level NVIDIA driver instead of the Nouveau DRM driver. Libdrm and related components are not used in the port.

The decision to use NVIDIA's code rather than the Nouveau driver was made based on ease of porting - the NVIDIA driver is designed to be portable and uses common code across Windows, Linux, FreeBSD and Solaris builds. The NVIDIA driver is also of higher quality and better supported.

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://discuss.haiku-os.org/t/haiku-nvidia-porting-nvidia-gpu-driver/16520

Google wants to reboot the kernel in the background:

21/03/2025

Engineers from Google have published a set of patches for discussion by Linux kernel developers with the implementation of the Live Update Orchestrator (LUO) subsystem, designed to update the kernel in Live mode. Unlike mechanisms such as livepatch, Ksplice, kpatch and kGraft, the new system is not limited to the ability to apply individual fixes to the running Linux kernel, but allows you to fully reboot and update the kernel without stopping the operation of individual devices. The project is based on a set of KHO (Kexec HandOver) patches to the kexec mechanism, used to load a new kernel from an

already running Linux kernel without a physical reboot.

The main application area of LUO is cloud environments, where it will be possible to update the KVM hypervisor without disrupting the operation of running virtual machines. In particular, it will be possible to suspend virtual machines while rebooting the kernel with the hypervisor, while keeping all devices attached to the virtual machines in working order.

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://lore.kernel.org/lkml/20250320024011.2995837-1-pasha.tatashin@soleen.com/

Raspberry Pi system image generator:

21/03/2025

The developers of the Raspberry Pi project have presented the rpi-image-gen toolkit, designed to simplify the generation of custom system images for Raspberry Pi boards. The advantages of the toolkit include: fast assembly due to the use of ready-made binary packages; the use of identical versions of libraries and applications with the Raspberry Pi OS; the ability to configure an arbitrary partition layout and use file system encryption; support for generating a list of used packages and checking for the absence of vulnerable versions of programs. The code is written in Shell and is distributed under the BSD license.

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://www.raspberrypi.com/news/introducing-rpi-image-gen-build-highly-customised-raspberry-pi-software-images/

ReactOS 0.4.15 Released:

22/03/2025

After more than three years of development, a new release of the ReactOS 0.4.15 operating system is presented, aimed at ensuring compatibility with Microsoft Windows programs and

drivers, and also offering a Windows-style design. An installation ISO image (117 MB) and a Live build (in a zip archive 85 MB) are available for download. The project code is distributed under the GPLv2 and LGPLv2 licenses.

"This release is a culmination of the work of numerous contributors since 0.4.14 was branched in 2020. This has been the largest release to date. There are nearly 8 times more commits going into this release than in 0.4.14. We are proud of the progress we have made, and are eager to continue with this growth"

https://reactos.org/project-news/reactos-0415-released/

NixOS offers a method to protect against backdoor substitution like XZ:

23/03/2025

A repeatable build mode is proposed for inclusion in the nixpkgs package repository used in the NixOS distribution, which allows identifying cases of backdoors being introduced into the code, reminiscent of the incident with the XZ project. The presented protection method allows detecting modifications in the archives with the source code of the release, which are absent from the repositories with the code.

The essence of the method is that the source code of the new version of the application is assembled twice - the first time from the code downloaded from the git repository, and the second time from the code distributed in ready-made archives. If the resulting binary files differ from each other, there is reason to suspect the presence of hidden modifications in the repository or in the archive file with the code.

https://github-com.translate.goog/NixOS/nixpkgs/pull/391569?_x_tr_sl=auto&_x_tr_tl=en&_x_tr_hl=en-US&_x_tr_pto=wapp

Unprivileged application isolation system:

23/03/2025

The Landrun project has begun developing a new system for isolated execution of individual applications. The isolation uses the Linux Landlock kernel LSM module, which allows you to do without executing privileged operations during the creation of a sandbox environment. The way it works, Landrun is close to the Firejail utility, but differs in a simpler implementation, lightweight, and the ability to work under a regular unprivileged user without delivery with the suid flag. The project code is written in Go and is distributed under the GPLv2 license.

The Landlock mechanism allows unprivileged programs to restrict the use of Linux kernel objects such as file hierarchies, network sockets, and ioctl. Unlike namespaces and system call filtering, the isolated environment is formed by the Linux kernel as an additional layer on top of the existing system access control mechanisms. To interact with the Landlock subsystem, the landrun utility uses the go-landlock library from the LandLock developers.

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://github.com/Zouuup/landrun

Release of Finnix 250:

23/03/2025

The release of the Finnix 250 Live distribution is presented, dedicated to the 25th anniversary of the project (the first version of Finnix was published on March 22, 2000). The distribution is based on Debian without a desktop environment and everything is done in the console and provides a selection of utilities for the needs of the system administrator. It includes more than 600 packages with all sorts of utilities. The size of the iso image is 528 MB.

https://translate.google.com/website?sl=auto&tl=en&hl=en-US&client=webapp&u=https://blog.finnix.org/2025/03/22/finnix-250-released/

ubuntu uk podcast

The Ubuntu Podcast covers all the latest news and issues facing Ubuntu Linux users and Free Software fans in general. The show appeals to the newest user and the oldest coder. Our discussions cover the development of Ubuntu but aren't overly technical. We are lucky enough to have some great guests on the show, telling us first hand about the latest exciting developments they are working on, in a way that we can all understand! We also talk about the Ubuntu community and what it gets up to.

The show is presented by members of the UK's Ubuntu Linux community. Because it is covered by the Ubuntu Code of Conduct it is suitable for all. The show is broadcast live every fortnight on a Tuesday evening (British time) and is available for download the following day. podcast.ubuntu-uk.org

RETURN TO CONTENTS

Command & Conquer:

Written by Eric

Last month, we looked a bit more in-depth at jobs, foreground and background processes. This allows me to segue into screen. Though we chatted about it before, (2022?), it was to segue into tmux. We can touch on tmux again - it is great - or tilix, but know that in the real world, you will probably connect to a jump host or a server that has only the OS installed, so things like tmux will not be available for you to use. Hence why I like to start off with screen, not that I'm old-skool or anything, I like tmux/tilix more, but trust me, you will get stuck with screen more often than not.

What is screen? Well, screen is a virtual terminal that allows you to run jobs, even after closing your terminal (continues the process). To demonstrate what I mean, open a terminal and type: evince

Your document viewer opens. Now close the terminal and you should get a warning, but close it anyway. What happened to evince? So are we on the same page? The evince process was terminated when your terminal emulator was closed. We don't want this when working on a remote machine. Imagine you have to zip up log files on remote machines and copy them elsewhere and you also need to go home as it is after five and you would rather not sit around for another hour at work, or you are busy hacking your school computer, running a brute-force attack, but it could take hours and you need to be home by 9, you don't want to stop dead in the water and start again from scratch tomorrow, you want it to finish.

What is nice about screen, is that you can have a few instances running at once and switch between the processes running in each. However, do not assume it is installed. You may need to add it with: sudo apt install screen – on both ends where you intend to use it. The man page for screen is very long, and it

uses that emacs type notation for key bindings, that look like this: C-a meaning CTRL + a. (Important, file for later use).

Fire up your terminal and type: screen

edd@gift () - byobu
GNU Screen version 4.09.01 (GNU) 20-Aug-23
Copyright (c) 2018-2023 Alexander Naumov, Amadeusz Slawinski
Copyright (c) 2015-2017 Juergen Weigert, Alexander Naumov, Amadeusz Slawinski
Copyright (c) 2010-2014 Juergen Weigert, Sadrul Habib Chowdhury
Copyright (c) 2008-2009 Juergen Weigert, Michael Schroeder, Micah Cowan,
Sadrul Habib Chowdhury
Copyright (c) 1993-2007 Juergen Weigert, Michael Schroeder
Copyright (c) 1987 Oliver Laumann
This program is free software; you can redistribute it and/or modify it under
the terms of the GNU General Public License as published by the Free Software
Foundation; either version 3, or (at your option) any later version.
This program is distributed in the hope that it will be useful, but WITHOUT
ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS
FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.
You should have received a copy of the GNU General Public License along with
[Press Space for next page; Return to end.]
u 24.04 0:-* 1d22h 1.21 20x2.0GHz 62.5G13% 1.8T42% 2025-01-08 12:21:56

If you are directly on a TTY, you should not see this, it would seem that nothing happened.

However, we can sort of get a grasp by doing the following.

Type:

[redacted]

press CTRL + a let go of the keys and press c

[redacted]

press CTRL + a let go of the keys and press c

[redacted]

press CTRL + a let go of the keys and press p

press CTRL + a let go of the keys and press n

What happened?

Now press

press CTRL + a let go of the keys and press 0

press CTRL + a let go of the keys and press 2

press CTRL + a let go of the keys and press “

You should be able to navigate with the up and down arrow key.

If you use it often, it will be good to know:

press CTRL + a let go of the keys and press A (Capital a) and name this screen something.

I renamed the middle one, “Centre” and when I bring up my list, I can see it. Your turn.

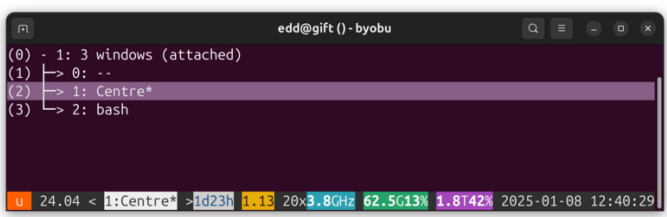

So now, instead of it just saying “bash”, it has a name and you should see a ‘*’ indicating which screen you are currently on. Start a ping in one of those screens and switch away and go back.

If you were wide awake, you may have noticed “detach” in the man page, obviously it is:

press CTRL + a let go of the keys and press d.

You should have your normal bash prompt back (some distributions will let you know with a message, like Ubuntu, but don’t count on it).

Now for the part I want you to compare with what happened to evince, close the terminal. Open a new one, it does not even have to be a gnome terminal.

Now type: screen -ls

That (detached) message, is the one we just disconnected from, and since it is “recent”, we can use: screen -R -to reattach to that session, or if you detached more than one or had multiple sessions open, like I did, you can use the PID displayed, to connect to it, but that is a lowercase “r”. (Check the man page for more “d” and “r” options...).

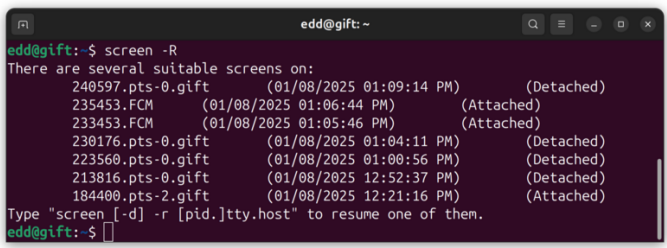
edd@gift: ~
edd@gift:~$ screen -R
There are several suitable screens on:
240597.pts-0.gift (01/08/2025 01:09:14 PM) (Detached)
235453.FCM (01/08/2025 01:06:44 PM) (Attached)
233453.FCM (01/08/2025 01:05:46 PM) (Attached)
230176.pts-0.gift (01/08/2025 01:04:11 PM) (Detached)
223560.pts-0.gift (01/08/2025 01:00:56 PM) (Detached)
213816.pts-0.gift (01/08/2025 12:52:37 PM) (Detached)
184400.pts-2.gift (01/08/2025 12:21:16 PM) (Attached)
Type "screen [-d] -r [pid.]tty.host" to resume one of them.
edd@gift:~$

Try it yourself quickly, I’ll wait....

One last thing: just like screens, you can also name your sessions (though I have never used it, as usually my session is one screen to one server), and you can bake that into a configuration file, but I feel we are stepping outside newbie territory. When starting a screen session, use -S (capital “S”) followed by a name (I named it FCM). You can also name a session from *inside screen, you just need to list all the open screens first, with:

press CTRL + a let go of the keys and press “ press CTRL + a let go of the keys and press :

You should now see a “:” at the bottom of the terminal. Type sessionname (one word)

Start naming and press < enter >

Now for the fun part, to kill a screen:

press CTRL + a let go of the keys and press k

and what about killall, I hear you ask, well:

press CTRL + a let go of the keys and press \

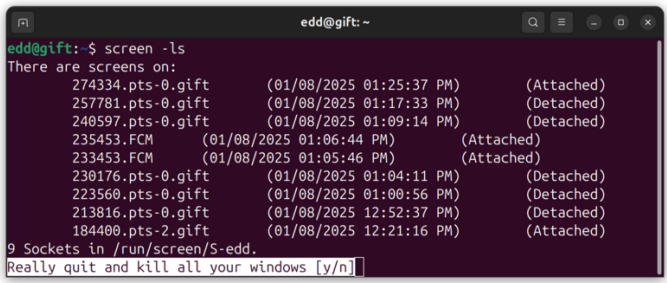
edd@gift: ~
edd@gift:~$ screen -ls
There are screens on:
274334.pts-0.gift (01/08/2025 01:25:37 PM) (Attached)
257781.pts-0.gift (01/08/2025 01:17:33 PM) (Detached)
240597.pts-0.gift (01/08/2025 01:09:14 PM) (Detached)
235453.FCM (01/08/2025 01:06:44 PM) (Attached)
233453.FCM (01/08/2025 01:05:46 PM) (Attached)
230176.pts-0.gift (01/08/2025 01:04:11 PM) (Detached)
223560.pts-0.gift (01/08/2025 01:00:56 PM) (Detached)
213816.pts-0.gift (01/08/2025 12:52:37 PM) (Detached)
184400.pts-2.gift (01/08/2025 12:21:16 PM) (Attached)
9 Sockets in /run/screen/S-edd.
Really quit and kill all your windows [y/n]

And that should cover the basics for any newbie, I really don't think you will need more in a loooong time, but have a look at the man page, and see that this is a grand daddy page! There are options to fit everyone's use case. I hope you followed along in your terminal, or it will be forgotten tomorrow.

Any mistakes, please write to: *misc@fullcirclemagazine.org*

***Erik** has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.*

RETURN TO CONTENTS

HowTo: Network Basics

Written by Erik

Chatting with a person in another group on Telegram, I mentioned that I wanted to set up a Pi-hole. The other person said that they would like to do so too, but they shared a house with three other people and the issue was that they had a VLAN set up to keep the users separate from each other, and the cameras separate. He did not know enough networking to add a Pi-hole in the mix. When I asked what his current DNS was, he told me it was 127.0.0.1... (localhost).

When I was knee-high to a grasshopper, I got a rubber-keyed ZX Spectrum computer for my birthday. Though I knew nothing about computers, what fascinated me was that one could move ‘sprites’ across the TV. It was either in the manual or in a magazine, that they showed you how, with BASIC. The catch was that you had to sketch out your sprite on graph paper, add up the rows of blocks, and use that number to ‘draw’ with. Challenge accepted! I got some graph paper, and went to work. The math was a bit strange, but I saw why immediately, as it allowed blocks in different columns to have a unique number. I went to town, making crazy sprites... Just to find out that the book taught you to move only an 8x8 block of ‘blocks’ and nothing else. All my awesome sprites were at least four to eight times as large! The frustration was real. I went on then to make 8x8 sprites and move them across the screen, by redrawing and deleting them in BASIC. It was slow and jarring and flickering, and a horrible experience overall. I hated it. As a kid, I could not understand why I could not make smooth scrolling sprites like I saw on the tapes that came with the computer, much less control the sprite to do what I wanted. It just flickered from left-to-right, or right-to-left as that was all I managed to change. I gave up on a bad job. Though my budding career as a games programmer was crushed right there, what stuck with me was the patterns of ‘blocks’ I made for the sprites and their values. The ‘blocks’ I coloured in on

the graph paper made this strange wave pattern, where they would fill up from the right-hand-side until they reached a number, then do it again to reach the next number. They were 1,2,4,8,16,32,64,128. So if you wanted to 'count' to 15, you would start at the one then add the two, then add the four, then add the eight, but now you would be stuck, so you would need to start the process over, keeping the eight 'coloured in', one, two, four, get stuck, leave it coloured in, for a total of twelve and start over until all the blocks were coloured in and you would have one less than the next block over, sixteen. I started to recognize these patterns and I could tell you the values of most of them without actually counting up the values. (Yes, I spent a lot of time making sprites).

What I was doing was actually binary maths, but I did not know it at the time. Actually a long time, as we did not cover binary in school. So it was at my third job, as a technician, when the boss got us some training. The part everyone had an issue with was networking. IP addressing, which I oddly recognized as sprite 'blocks', had it all come flooding back. I did not know any theory behind it, to me it was all a 'sprite', so adding the ones and zeroes made perfect sense.

128	64	32	16	8	4	2	1	Total
								0
								0
			■					16
		■	■	■				56
			■					16
								0
								0

In my example, the 'sprite' for the plus-sign would be (0.0.0.16.56.16.0.0), and my network address is (10.0.0.2.255.0.0.0) - do you see where I make the connection? IP addresses make for some ugly sprites, but you can draw them out and start to recognize the patterns.

Then they had us break apart those octets to binary, useful for network engineers, not really for home users, but file it away somewhere. However, the binary was just another 'sprite' to me. Those 'patterns' I learned from making 'sprites' also came

in handy again, as I knew, say, the last two blocks would total to 192 if I needed them coloured in. So if we needed to work an IP address backwards (the part that seemed to break my fellow technicians' brains), say the number 210 for the first number (and say I could not remember the pattern). I would know that the first three blocks together would be 224 and 224 is larger than 210, so it would have to be the last two blocks, value as stated above, and I can immediately put ones in the last two blocks, do the subtraction, end up with 18, and put a one in the 16's column and one in the 2's column, then zeroes in the rest. Really, that's how easy it was. If this way of thinking about it helps anyone, I'll count it as 'mission accomplished'.

At this point you may be saying, well we only use DHCP at home and at work, so what's the point? Well you may encounter it somewhere, maybe when you start setting up your smart IOT devices, or faffing with your router making subnets that can access the internet and some that cannot, or you end up sharing a house where someone else set it up and you need to add a Pi-hole.

Let's break an IP address down, take it apart, if you wish. If we look at an IP address of say 192.168.1.2/24, like a lot of home routers give you, the first three numbers, 192.168.1 – is the network part, and the last number, the 2, is the host portion. I'm just looking at IPv4 here, it's easier to grasp. The octet in IPv4 is always 32 bits long (8x8 on graph paper). There is also a subnet mask that, you guessed it, masks the network portion of our IP. The way it does this is it fills in (masks) the network portion with all ones, and leaves all zeroes for the host portion. Thus in our example above the first three numbers will get 11111111.11111111.111111.00000000 and the last one zeroes. That translates to 255.255.255.0. The next logical column, that would follow 128, would be 256. So if you fill in all the blocks to just before 256, the total is 255. You can do the maths if you want to, that will be 127 (all the blocks before 128) plus 128, I won't lead you down the wrong path. That said, the subnet mask can move, so it will not always be 255.255.255.0. If I were to mask, say, the first 20 bits, that

would translate to 255.255.240.0 and that really is the fundamentals of an IP address. There are two IP addresses in each range you cannot assign to a device from the host portion, that being the first and the last, 0 and 255. Other than that, go bananas (OK, 127.0.0.1 is reserved as the loopback address, but other than that... LOL).

Now, when you deal with your ISP and they give you a fixed IP, this will mostly happen for businesses and not home users, they may slice that pie smaller than 255.255.255.0. Then like before, the first and last IP addresses in that range cannot be used, so just looking at an IP address you can't tell unless you are salted, you need to check if the host portion translates to all 0's or all 1's. The company I used to work for used these fixed IP addresses to do whitelisting on our infrastructure. If you see an IP address 10.128.224.64, it will translate to: 00001010.10000000.11100000.01000000 (and usually we lose the full stops). If my subnet mask is 255.255.255.0, that^^^ is a host address, but if my subnet mask is 255.255.255.224, it is not any more. How do I know? Well, if I mask out 1's up to 5 characters from the end, I'm left with five 0's, meaning it is a network address, the first address I told you that you cannot use. So as long as it is not all 0's or all 1's you can assign it to a device. So do not fool yourself into thinking you can tell just by looking.

There is another caveat. There are private IP addresses that cannot be routed. Most of the time, your home router will hand out IPs on the 192.168.0.x range that usually has a network mask of 255.255.255.0, meaning you can have 255-2=253 number of addresses and this is what most people see. I just want to point out that there are three sets of IP addresses that are private. 192.168.0.0-192.168.255.255, 172.16.0.0-172.31.255.255, and 10.0.0.0-10.255.255.255

In Ubuntu Gnome, we see the networking section under the settings, right at the top in the left column. By default, we cannot change a network unless we try to connect to it.

I'll use WiFi as an example. If you look at the image above, you will notice that I can only click on any one of these. The magic happens only after you do.

Now that I have clicked on a few, I suddenly see a gear icon displayed.

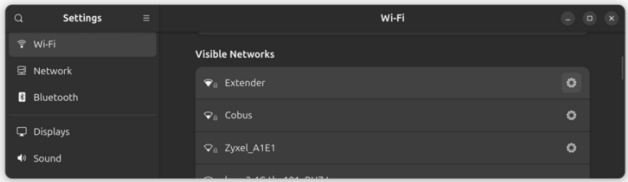

If you want to follow along, but you do not have wireless, you can use the "network" button below "Wi-Fi" if you have a cable, or if you have neither, you can do what I did and click on any of the detected SSIDs. We just want the configuration window to display, we don't need to connect to the foreign networks. Here it is:

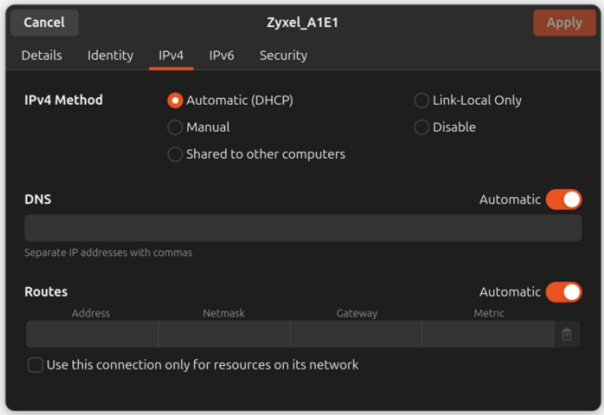

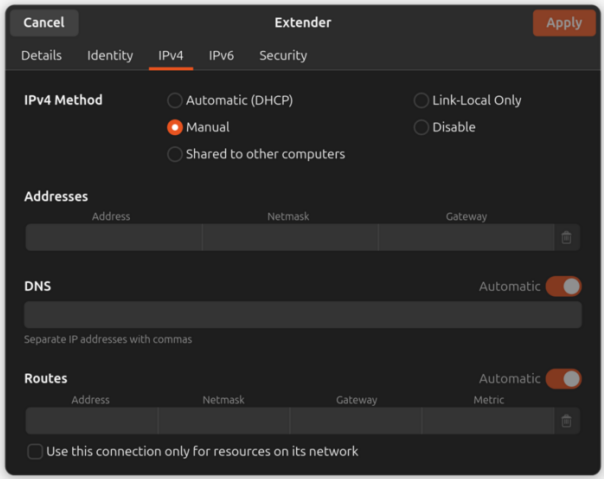

Notice the tabs along the top row. The flower will open if you change that IPv4 method from DHCP to manual.

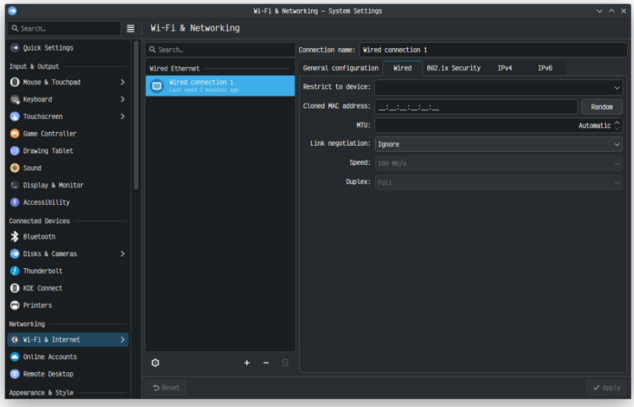

On KDE, one can go directly to “configure network connections” or use system settings, but in that case, the networking section is near the bottom.

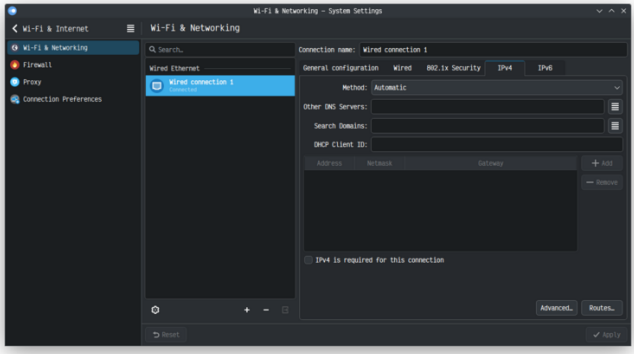

In the next issue we can discuss topics like NAT, network address translation, and maybe how a router works.

As always, complaints to: *misc@fullcirclemagazine.org*

***Erik** has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and*

oscilloscopes, laying cable, to scaling 3G towers, he's done it.

RETURN TO CONTENTS

HowTo: Trading up to Linux – Part 3

Written by Alan German

We have now arrived at that scary moment when it's time to install Linux onto our computer's hard drive. However, there is absolutely no need to be anxious about this process. I have performed such installations many times without incident. Furthermore, our method for installing Linux will retain the option to use Windows, and is easily reversible if at any point in the future you decide that Linux isn't right for you after all. We will go through the installation process step-by-step. and for those who may need even more comprehensive instructions, a complete, fully-illustrated guide is available at *https://linuxnorth.wordpress.com/installing-linux-mint-as-a-dual-boot-system/*

Notwithstanding the above reassurance of the ease and safety of installing Linux, we will be making significant changes to the hard drive, notably by modifying the partition structure. Consequently, our first step should be to make a full disk image backup of our system so that we can restore the disk to our starting point in the unlikely event that the installation process goes south. In any case, it's always a good idea to have a recent system backup so the process of installing Linux provides a good excuse for doing so now.

If you are not employing a specific backup strategy and want to use a conventional Windows-based backup system, try EaseUS Todo Backup (see: Finding a Backup Program *https://opcug.ca/Reviews/BackupPrograms.htm*.) An alternative, since we now have some experience using live-USBs and the Linux OS, might be Foxclone (see: Exploring Linux – Part 33 – A Backup Solution for Linux Systems. *https://opcug.ca/Reviews/Foxclone.pdf*). Either of these programs will readily create a disk image of the entire system as a backup.

So, with our backup in hand, our first step in the installation process is to reboot the computer into the Linux desktop using our Linux Mint live-USB. But, this time, instead of exploring the menu options and trying the available applications, we will focus our attention on the top-left corner of the desktop where we will find an icon labelled 'Install Linux Mint'. Not surprisingly, we are going to double-click on this icon and launch the installation process.

What follows is a series of screens in the form of an installation wizard that will guide us through the process. Most of the inputs required by the dialogue boxes will be self-explanatory. Where this might not be the case, you should find detailed explanations and screenshots in the following text. So, take a deep breath, and let's get started.

The very first screen is straightforward. It is seeking the language that is to be used by the wizard. The default selection is English, and to accept this option it is merely necessary to press the Continue button in the lower-right corner of the window. The next screen is similarly easy to deal with. It presents the default option for the keyboard layout as English (US) which is the most likely scenario for users across North America. So, once again, we can simply press Continue.

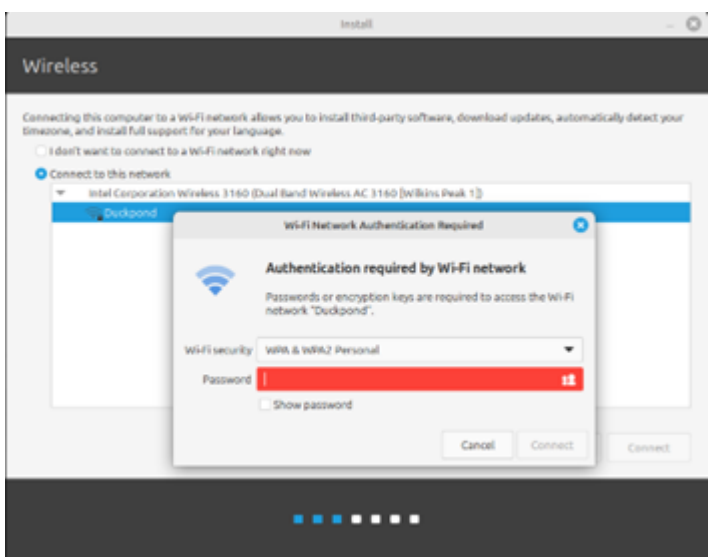
Install
Wireless
Connecting this computer to a Wi-Fi network allows you to install third-party software, download updates, automatically detect your timezone, and install full support for your language.
I don't want to connect to a Wi-Fi network right now
Connect to this network
Intel Corporation Wireless 3160 (Dual Band Wireless AC 3160 [Wilkins Peak 1])
Duckpond
Wi-Fi Network Authentication Required
Authentication required by Wi-Fi network
Passwords or encryption keys are required to access the Wi-Fi network "Duckpond".
Wi-Fi security
WPA & WPA2 Personal
Password
Show password
Cancel
Connect

Figure 1. Wi-Fi setup

Now, we come to our first real choice. The screen entitled Wireless (Figure 1) gives us an option to connect to our local W-Fi network. My strong recommendation is to select the entry for your specific network (as I have done for Duckpond). Click on the radio button labelled Connect to this network, and press the button labelled Connect. A pop-up dialogue box now requests your Wi-Fi password and activates a second Connect button. Pressing this button returns us to the initial screen where we can finally press the Continue button. This process establishes Wi-Fi connectivity that can be accessed as part of the installation process but, more importantly, provides the information that enables the installed version of Linux to automatically connect to the local Wi-Fi network on each subsequent bootup.

Now that we have established Wi-Fi access, we can select the option on the next screen to Install multimedia codecs. This will allow our Linux system to play audio and video files that require encoding systems that are not in open formats. Most of us will not have any reservations about using software that is not solely open-source (we have been Windows users after all!) and will be happy to simply press Continue.

Now, an information screen pops up that asks if we want the Linux installer to try to unmount disk partitions that are in use and will probably indicate that the disk is something like /dev/sdc. It won't be obvious to those currently unfamiliar with Linux that this is actually referring to our live-USB drive. On my system, the devices (dev) in use are sda (my main hard drive), sdb (an SDcard used for backup), and sdc (the USB flash drive with the Linux software). The displayed text is noting that, if the USB drive is left mounted, we may be able to create additional partitions on the drive. Since we have no need to do this, the simple option is to press the Yes button and have the installer unmount the drive. The software required to complete the installation has already been read into memory so we have no need for the live-USB drive to remain mounted.

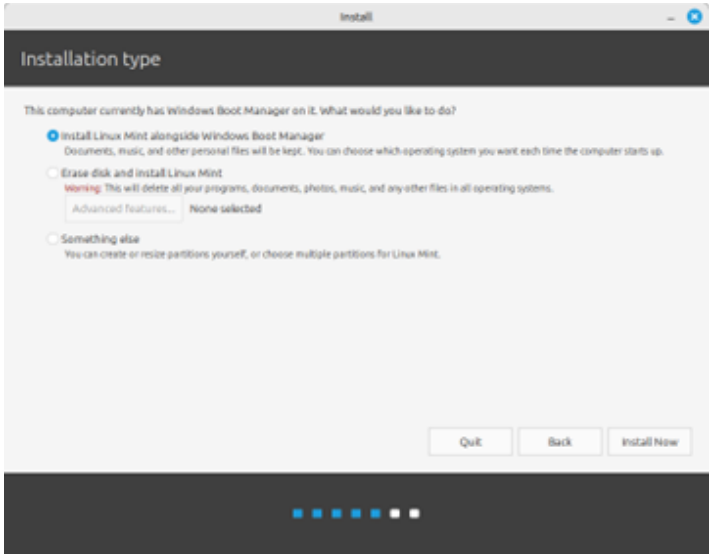
Install
Installation type
This computer currently has Windows Boot Manager on it. What would you like to do?
Install Linux Mint alongside Windows Boot Manager
Documents, music, and other personal files will be kept. You can choose which operating system you want each time the computer starts up.
Erase disk and install Linux Mint
Warning: This will delete all your programs, documents, photos, music, and any other files in all operating systems.
Advanced features...
None selected
Something else
You can create or resize partitions yourself, or choose multiple partitions for Linux Mint.
Quit
Back
Install Now

Figure 2. Set the installation type

The next screen (Figure 2) determines how Linux will be installed on our hard drive. The default entry is to Install Linux Mint alongside Windows Boot Manager. This is the option that we wish to use as it will automatically shrink the Windows partition, create a new partition in the free space, and install Linux in this new partition.

So, press the Install Now button. This will result in a new screen (Figure 3) with a graphical representation of how the Windows and Linux disk partitions are going to be sized. The installer will make a choice of how to split the existing Windows partition but this may be changed by simply sliding the divider between the two partitions, to the left to create a larger Linux partition, or to the right to maintain more space for Windows.

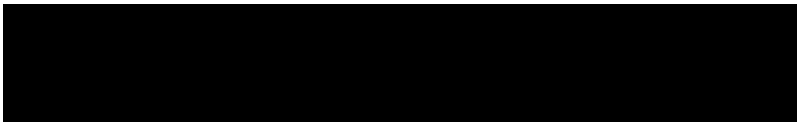

https://linuxnorth.wordpress.com/manual-partitioning-revisited/

Linux Mint will run easily in a 20 GB partition and, in my case, the installer has chosen to make the new partition almost twice as large at 38 GB. Note also that the Windows partition will be shrunk to about 210 GB. The used space on this partition is currently about 176 GB so there will still be 34 GB of free space available for use in Windows. As a first option this all sounds quite reasonable and, in any case, the partition sizes can always be adjusted should this become necessary at some point in the future.

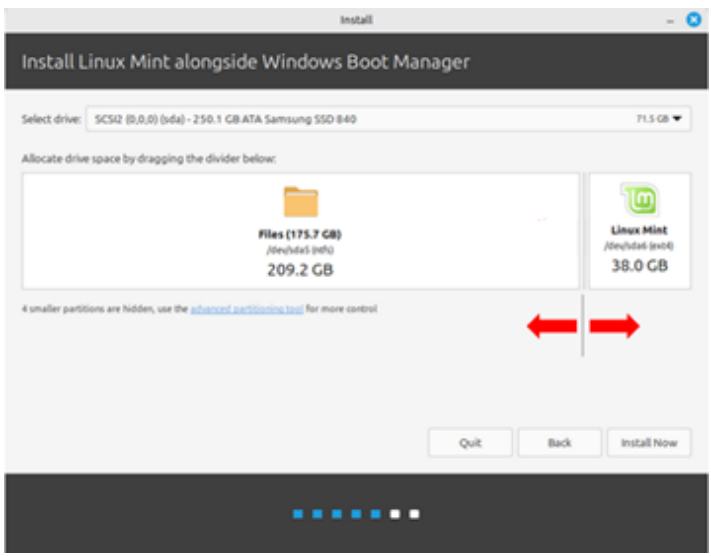

Figure 3. Establish the partition sizes for a dual-boot system

Next, we see two further warnings. Firstly, the installer reminds us that the changes to the disk partitions must be made and that this cannot be undone. Clearly we need to allow the new partition scheme to be established in order to create our dual-

boot, Windows-Linux system. This process cannot be undone by the installer but we have our disk image backup that we can use to restore our original setup should anything go wrong (which it won't!). Pressing the Continue button pops-up a second reminder informing us that changes will be written to the disk. In this case, the bottom line will be something to the effect that "The following partitions are going to be formatted: partition#6 of SCSI2 (0,0,0)(sda) as ext4". This is simply telling us that the new Linux partition (Partition No. 6) on our hard drive (designated as sda) will be formatted using the Linux ext4 file system. Once again, this is something that needs to happen, so press Continue to allow the changes to the partitions to take place.

A map of the world is now displayed asking Where are you? The highlighted zone is likely to be labelled as a city near you and there is normally no need to change this location. For example, even though I am located in Ottawa, my computer displays Toronto (and there is no option to select Ottawa!) The installer doesn't want to know the actual geographical location; it's just going to use this to set the appropriate time zone. The highlighted area of the world is the Eastern Time Zone which, for me, is correct. So, if your situation is similar, no changes are necessary, and you can simply press Continue.

The next screen (Figure 4), labelled Who are you?, has several blank text boxes that we need to complete. Enter your name (first name, full name, or a pseudonym is perfectly fine), and the installer will populate the next field with a suggested computer name. The entries to all of the data fields are totally flexible so make these names whatever you prefer.

Next, we need to enter a username and password to use as our login credentials. Note that the username must begin with a lower-case letter, otherwise an error message will pop up indicating "Must start with a lower-case letter." Choose a username and password that make sense to you, and make a note of the parameters selected. The password, in particular, will be required both for logging in to Linux and also for any command that needs authorization (in much the same way as

required by the Account User Control system in Windows).

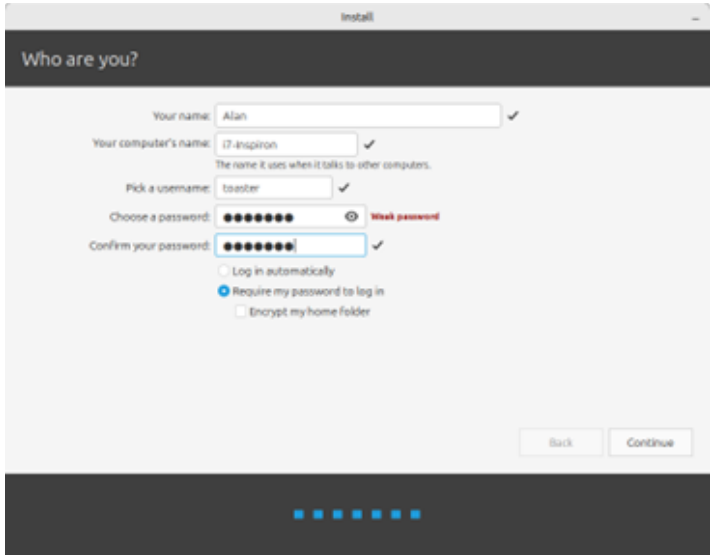

Figure 4. User Configuration

The final options on this screen are radio buttons for the login method and a checkbox to encrypt the home folder in the Linux filesystem. My recommendation, and my practice, is to accept the default entry to Require my password to login and to leave the box labelled Encrypt my home folder unchecked.

That completes the data entry process. Press the final Continue button and a series of information screens is displayed as the installation proceeds. There are also a series of messages and a progress bar displayed along the bottom of the window. You can scroll through the information screens, using the left and right arrows at the edges of the window, or simply let the installer run to completion.

On my computer, this part of the installation process took about ten minutes. I was then presented with a pop-up dialogue box indicating that the installation had finished with control buttons to either Continue Testing or Restart Now.

We have already used the live-USB to test Linux on our computer, so now it's time to take the final step and press

Restart Now. The computer will reboot, but now will display the GRUB (GRand Unified Bootloader) menu, essentially offering the choice to boot into either Linux or Windows. The default boot sequence is Linux; however, there is a 10-second delay before the computer boots into the default OS. At any time during the countdown, the arrow keys can be used to select another menu option. The other option of choice is likely to be to boot into Windows which, for my computer, means selecting the option for Windows Boot Manager (on /dev/sda1).

We have now installed Linux in a new partition on the hard drive of our computer, while maintaining the Windows operating system and its associated files in the original partition albeit with a somewhat reduced size. The grub menu allows us to easily boot into either of the two operating systems so we have successfully created a dual-boot environment. The next step will be to update the Linux software and start to customize the OS to suit our requirements. But, that will have to wait for the next installment in the present series of articles.

Alan is a computer enthusiast based in the Great White North where he is an active member of the Ottawa PC Users' Group (*https://opcug.ca*) and maintains the LinuxNorth blog at *https://linuxnorth.wordpress.com*

***Alan** is a computer enthusiast based in the Great White North where he is an active member of the Ottawa PC Users' Group (https://opcug.ca) and maintains the LinuxNorth blog at https://linuxnorth.wordpress.com*

RETURN TO CONTENTS

HowTo: Latex - Babel

Written by Robert Boardman

Welcome back to Full Circle and to Latex. This time, I am going to explain a few ways to install new packages to Latex on Linux. My “development” machine runs Mint 22.1, which is a Ubuntu derivative. This information is to alert you to the possibility the locations of my files on my machine may not match the locations of the same files on your machine.

In last month’s column, I reviewed the history of Tex and indicated why it was written using packages. In addition to the very limited computing resources available when it was developed, the package structure points out most people who write have specialized needs. They are primarily mathematicians or chemists or chefs or translators or whatever. Most people do not need or use all the tools available in Latex or word processors. Most use a set of the same features in almost every document. With Tex / Latex, it is easy to make a template that includes all of the required packages and then generate the same kind of document every time. For example, if you never use graphs in documents then do not fill up computer memory with software that generates graphs.

There are a few thousand packages described in the two volumes of The Latex Companion (ISBN 978-0-13-465894-0). There are many more packages available at ctan.org. You should be able to find at least one package that will help you get the result you want. First you need to know if a desired package is already installed.

I use a tool to help me find files. The command is “locate”. If I remember correctly it is not part of recent standard Linux installations. If it does not exist on your system, it is a very small bit of code, easy to install from a terminal using apt or whatever method your operating system uses. It searches a database of file names so, before you use it the first time, type

“sudo updatedb”. The locate database should be updated regularly. It is a good candidate for the root cron. Once the database is built or updated you can use locate to find the package you want to use.

Another way to find out if the package you need has been installed is to search for it using tlmgr (TexLive manager). You need to know the name of the package so do a search at CTAN to find a package that fits your needs. On the command line, type tlmgr info package-name. It takes a few seconds and then generates several lines of information about the package. You could also use locate or find to do the same task.

If you have use of tlmgr, it will normally do the work necessary to install any extra package you need. It is similar to apt (or any standard system upgrading and installation program). An easy way to check if tlmgr is installed and/or usable is to run tlmgr update –list. It will go through the installed packages on your device and list all those which need to be updated. Since it checks your complete TexLive/Latex installation, this is not an instantaneous action. It takes a few seconds before you will see results. Of course if it is not installed or you do not have rights to use it, then you should see the error message quickly. The help available for tlmgr is extensive. When I imported it into OpenOffice Writer it spanned thirty-five pages.

Once you have done a couple of searches for Latex files you will see the general structure of the tex directory tree. If you record the location of the major file groups, you may decide you do not need to use locate. I use locate because it is easy for me.

On my Mint system, Latex packages tend to get installed in one of three places, either in /usr/local/texlive/2024/bin/x86_64-linux or in /usr/share/texlive/texmf-dist/tex/latex or in /usr/local/texlive/2024/texmf-dist/tex/latex. As you might guess, “2024” indicates the release year of texlive/Linux I use. On my system /usr/local/texlive/2024 is essentially the root of the tex directory tree.

What happens if you use the usepackage {name of desired

package} command, and the package is not installed? That depends on what interface you use to write Latex/Tex files. If I write Latex files with a text editor then I find out about a missing package during the compile phase, when the Tex file is converted to a PDF file. If a package is missing then the Tex file will not generate a PDF file. Instead an error will be generated saying there is a missing package.

If I use TexStudio (which I usually do), then I can see immediately if a new package is required. When I type \usepackage{ TexStudio shows me a list of all the packages available on my system. Obviously if I want to use a package not on the list I must install it before the compile will be error-free. Then I use tlmgr in a terminal to install it:

Sometimes the package on ctan.org is not structured in a way that tlmgr can use. This seems to happen most frequently with old packages. That is when the locate utility I mentioned becomes useful. Most packages add one (occasionally more than one) style file (.sty) to the tex directories. Some packages add a class file (.cls) instead. Here is the process I have used to install old packages that will not install with tlmgr:

1. *Download the package zip from ctan.org.*
2. *Unpack the zip file and read all the documentation if it exists. This will tell me two things: 1) will the package do what I need and 2) is there anything special or complicated about the installation.*

There is a WikiBooks page entitled Latex_Installing_Extra_Packages from 2022 September. It says Latex packages are installed in subdirectories in texmf/tex/latex named after each package. If in doubt about where packages are installed find “texmf/tex/latex” (or texmf/tex-dist/latex) on your system.

Some old packages require the first four steps in the process listed in WikiBooks and shown below. I avoid those packages if

possible. Other old packages will contain some short documentation plus either a .sty file or a .cls file. Then I follow the steps below:

1. *I use locate to find the installed .sty / .cls files, or I look for the piece of paper that I wrote this location on.*
2. *I make a new directory in the part of the tree that has the other .sty /. cls directories. I name this new directory the name of the package.*
3. *I copy the .sty / .cls file into the new directory.*
4. *I run texhash to rebuild the Tex directory database.*

The instructions from WikiBooks (2022) for installing new packages involve these steps:

1. *Extract the files. Run Latex on the supplied ins file. (See Comments below.)*
2. *Generate the documentation. Run Latex on the .dtx file. This will make a dvi file or a pdf if you have pdfLatex installed. (See Comments below.)*
3. *Install the files to a local directory in the ~/texmf part of the tree (probably ~/texmf/tex/latex/local). Using a local directory prevents any of the new package files from writing over the standard TexLive installation. To be consistent, name the directory for the new package the same as the package. Usually there is only a ty file to move into that new directory.*
4. *Update the Latex index of files. Most Important and easy to do. Simply run the command texhash and wait for it to complete. If you forget to run texhash, Latex will not know where the new file or files are and none of them will be available to use.*
5. *If the installation involved fonts, you also need to update the font map. Fonts can be installed in the system-wide directory and/or in the local (personal) directory. There are two different ways to update the Latex font map. (The newfont.map file will be provided with the font package.)*

Local: I direct you to *www.tug.org/fotts/fontinstall-personal.html* for detailed instructions. TUG discourages users from installing fonts locally. The process is more complicated and forces the use of updmap-sys every time after updmap-user is run in order for all fonts on the device to be recognized.

Whether installed locally or system-wide, the system file name database has to be updated after fonts are installed. You can use mktexslr (or texhash) to do that. If you installed Latex as root (not recommended), then you must use sudo -H mktexlsr.

Comments on the Four / Five Step process above.

The Wikibooks steps are unnecessarily complicated. Perhaps they were necessary in the past. For most packages there is a simple two-step installation process.

1. *In a terminal tlmgr install name-of-package*
2. *exhash*

If you use tlmgr update - -all the mktexslr command mentioned with fonts runs automatically once the update is finished. That is very convenient. I also run sudo texhash after the update just to be certain.

This has been sufficient for almost all the packages I have installed while writing these columns. Unlike the Wikibooks 5 Steps, I have used this two step process for font packages with no problems.

If you wish to (or need to) use the manual process instead of using tlmgr, then save yourself time and energy and read the documentation for the package. Usually packages built in the last ten years are complete with PDF documentation, and only require Step 3 and Step 4. It is unlikely you will need to do Step 1 or Step 2. It is unlikely you will need to generate any files or build a documentation file. This is generally true for all reasonably recent packages at CTAN. However, for some old packages this may be necessary. If you find a package on CTAN that requires Step 1 and Step 2, I suggest you search for a more up-to-date package that does what you need.

That is enough for this column. I hope you try to add a few new packages in the near future. Let me know if you have questions or comments. I will talk to you next time about more packages from ctan.

RETURN TO CONTENTS

KILOBYTE
MAGAZINE
Kilobyte Magazine is a fanzine for 8bit enthusiasts. It covers consoles, computers, handhelds and more, as well as new games for old systems. If you grew up with Commodore, Atari, Sinclair or Amstrad, this magazine is for you.
https://retro.wtf/kilobytemagazine/

HowTo: Drawing with Inkscape - Part 155

Written by Mark Crutch

Over a decade ago, way back in 2014, one of these articles covered the use of clipping paths to efficiently use several parts of a single bitmap image in a design, without the overhead of including multiple copies of the file data (part 32, FCM #92). More recently the same idea was revisited (part 148, FCM #208) with the introduction of the "Crop Image to Clip" option in the context menu for bitmap images. In this case, it's not so much that the image isn't included more than once, but rather that most of the unused parts are thrown away, resulting in separate embedded images that may require less storage space in the SVG file. This feature was added in Inkscape 1.3, but it was far from being the headline addition for that version.

The most lauded change in Inkscape 1.3 was probably the addition of the Shape Builder tool, which I covered in part 137 of this series (FCM #197). If you're not familiar with this tool, then I recommend reading that article before continuing with this one. As I described it there, "the Shape Builder can be thought of as a convenient way to perform certain Boolean operations between objects," which hints at its biggest limitation: it only works with vector objects. Until now.

As of version 1.4, it works with bitmap images too. Which gives us yet another way to chop up La Giaconda into smaller parts that can be freely rearranged. As before, I plan to isolate her head and hands, so the first step is to drop my bitmap image into Inkscape and draw some shapes marking the areas of the picture that I want to keep.

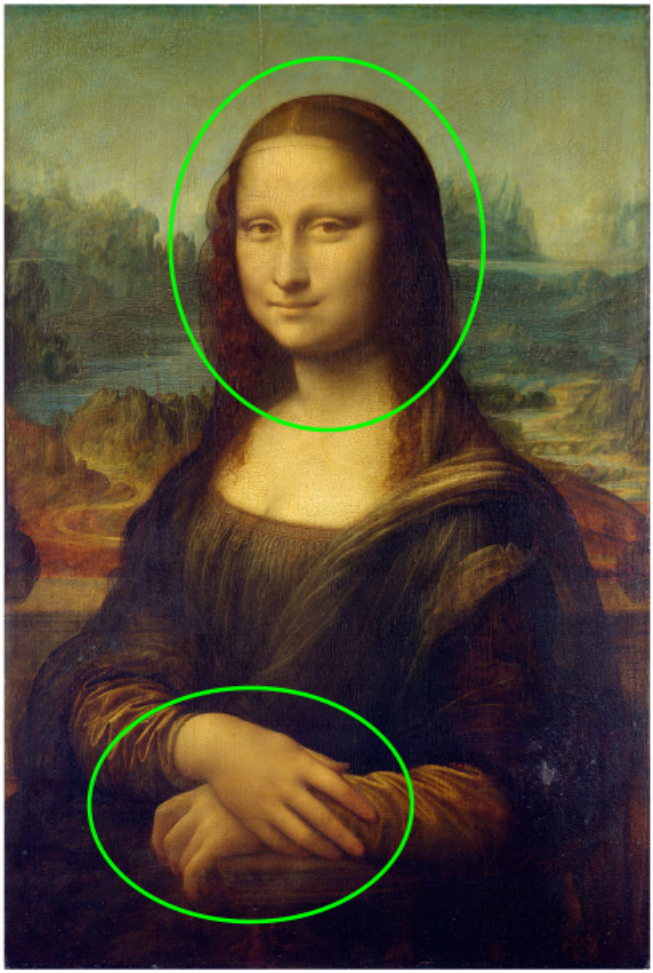

In the 2014 version of this technique I'd now be faced with grouping the image, then entering the group to clone it, copy the clone to the clipboard, exit the group, paste the clone, and use the two ellipses as clipping paths. It does the job well, but requires several steps and too much understanding about the structure you're ultimately making in the SVG file.

The 2024 approach, with Inkscape 1.3 or later, is a little simpler. Rather than creating complex arrangements of groups and clones, you just duplicate the image (Ctrl-D) until you have as many copies as you have shapes to clip with. Each copy is

then clipped with one of the shapes – so in this case I will duplicate once, and clip twice. Selecting one of those clipped versions shows a bounding box that makes it clear that the entire image is still present – so all you've done so far is drastically increase the size of your file, due to that duplication.

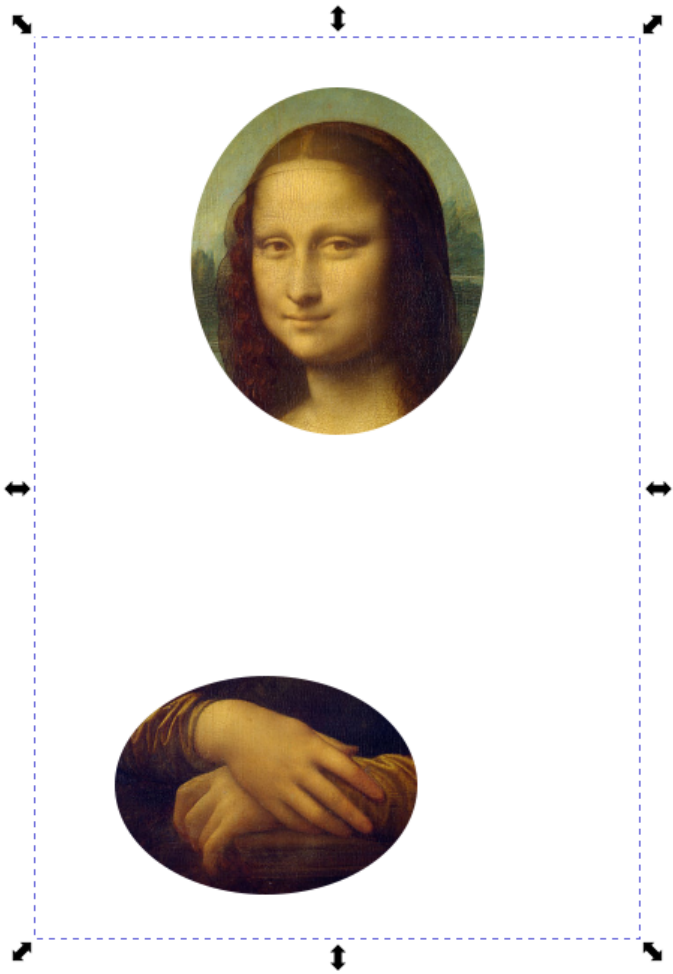

The last step, however, is to right-click on each clipped image and select "Crop Image to Clip". This throws away much of the image content, such that the bounding boxes are now only as large as they need to be, and the status bar will show that the image dimensions have reduced.

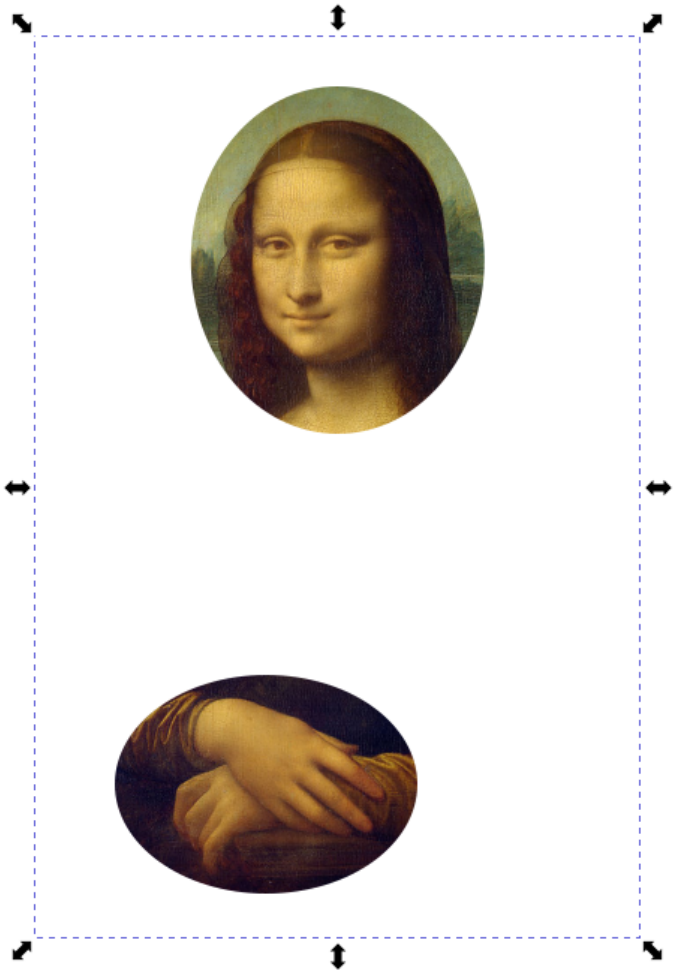

But this is 2025 and Inkscape 1.4 gives us a new approach that is simpler to apply than both of these. The result will be conceptually closer to the first, but the simplicity alone probably makes this the preferred approach in most cases. The steps are as follows:

Prepare the bitmap image and clipping shapes as before.

Select all of the objects.

Switch to the Shape Builder tool (default keyboard shortcut: X).

Click within each of the clipping shapes so that they are shown with a blue translucent fill (see image below).

Press Enter, or click on the Accept button in the tool control bar.

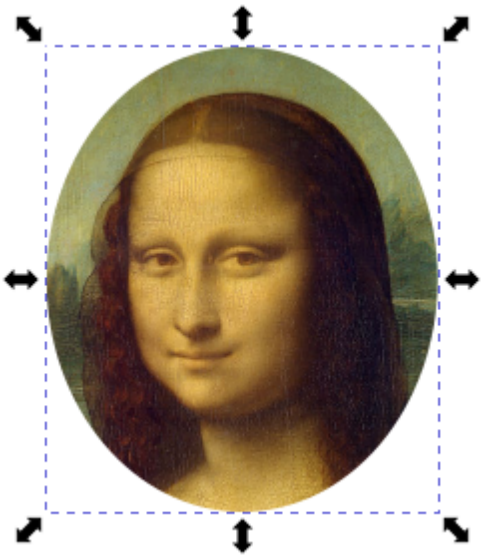

What you're left with is similar to the 2014 approach – clipped clones of the original image. But whereas the older approach left the original on the canvas (inside a group), this method moves the original bitmap data into the <defs> section of the SVG content – basically a part of the file that is used to store resources that are required to draw your design (such as gradient and filter definitions), but which do not directly appear on the canvas.

Because the original image is no longer present on the canvas, you can't select it in order to clone it again should you need to create another clipped shape – perhaps focusing on the eyes or the mouth, for example. At first this might seem like a bit of a limitation of this approach compared with the older method, but there is a very simple workaround that renders it a non-issue.

If you require another clone of the image, just duplicate (or copy and paste) one of your existing clipped shapes, then release the clip either via the context menu or the Object > Clip > Release Clip menu entry. This will leave you with a clone of the image from the <defs> section, plus a copy of your clipping path, which you can then delete or edit as necessary. Don't be surprised if the clipping path looks a little different to your original – Inkscape throws away the style data when the path is used for clipping in this way, as it's no longer necessary. What you'll see, therefore, is the default SVG rendering of a path with unset fill and stroke – which appears as though it's got a black fill.

Of course you're free to set it back to whatever fill and stroke you prefer for your clipping paths (that'd be transparent fill and bright green stroke, for me) if you want to reuse the path rather than deleting it. Note that if you clip "normally" rather than using the Shape Builder, Inkscape preserves the style information for the path. It's only the path shape data that is relevant for clipping, so whether the style of a clipping path is conserved or not is irrelevant for your final design.

In the event that you change your mind entirely you can, of course, undo the Shape Builder operation to return to the original situation of having the image directly on the canvas and no longer in <defs>. But what if you've saved and reloaded, such that you can't undo back to that step any more, or you've made other edits since then that you don't want to undo? Releasing one of your clipped shapes will, as above, get you a clone of the image plus the clipping path. Selecting the clone and using the Edit > Clone > Unlink Clone menu entry will turn the clone back into a real image on the canvas. But beware! This does not remove the image from the <defs> section, so you now have two copies of the image data in your SVG file, bloating its size. Assuming nothing else is using the <defs> version (i.e. you've removed all the clones from your document), you can use File > Clean Up Document to clear out any unused entries in the &ly;defs> section, including the image data.

When using the Shape Builder with a bitmap image, it's also extremely easy to create an 'inverse clip'. We've looked at how to clip the image to leave only the areas enclosed by paths, but it's just as easy to remove only those areas instead. Consider, for example, creating a version of the image with the face punched out (so to speak). For this, the steps are very similar to the earlier ones:

1. *Prepare the bitmap image and clipping shape.*
2. *Select all of the objects.*
3. *Switch to the Shape Builder tool (default keyboard shortcut: X).*
4. *Click on the image, outside of the clipping shape.*
5. *Press Enter, or click on the Accept button in the tool control bar.*

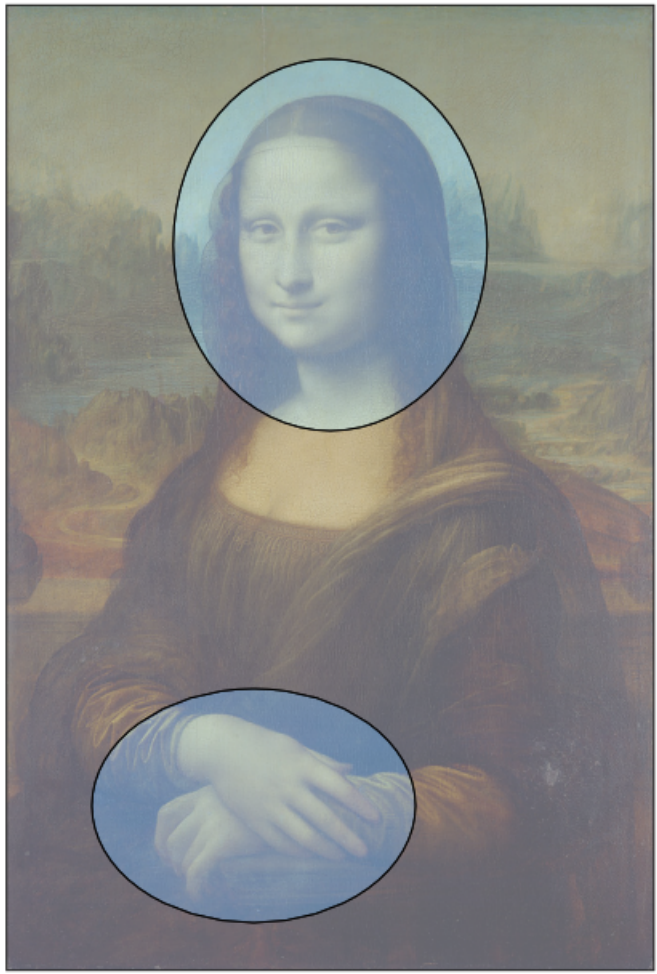

Perhaps what's most interesting here is the clipping path that's being used. It's no longer the path you originally drew, but rather a complex path which covers the image and has a hole in it that corresponds to your original clipping path. In other words, an 'inverse' clipping path. Release the clip to see it in all its unset fill glory.

Creating an inverse clipping path has historically required manual work with Boolean operations. More recently, Inkscape has had a direct means to achieve this via Object > Clip > Set Inverse Clip (LPE) but, as the name suggests, this uses a Live Path Effect rather than just creating a suitable inverse path. It also doesn't work on bitmap images, so it's a non-starter in this specific case. The Shape Builder approach is much, much simpler than trying to create a suitable inverse path yourself, so it's definitely the approach I will use in future.

If you do want your original, non-inverse, path back when you

release the clip, just select the inverse clipping path, use Path > Break Apart, and remove the outer part.

Although this specific example results in a complex clipping mask, that won't always be the case. For disjointed areas, the Shape Builder will more commonly produce separate clipping paths – as in the case of the separate ellipses for face and hands at the start of this article. Consider these two uses of the Shape Builder:

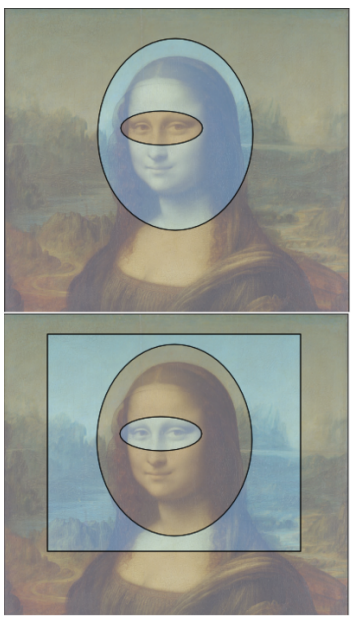

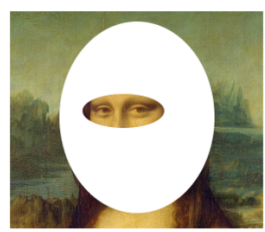

The top design results in a single complex clipping path consisting of an inner and an outer loop. The bottom design results in two clipping paths: one simple path for the eyes, and a complex path for the outer section. Each of the resultant shapes in the bottom design can be moved independently, which may not necessarily be what you want.

If you do need to keep several parts relatively positioned after clipping, I strongly advise grouping them immediately after the Shape Builder has done its thing. It is possible to unclip them all, use Path > Combine to turn all the clipping paths into a single complex path, remove the excess clones and then re-clip

a single remaining clone to create the same visible result with just a single clipping path – but there are very few cases where such extra effort would be warranted.

Will the Shape Builder be the last word in techniques for chopping up bitmap images into smaller parts? Only time will tell. But it's definitely easier and more intuitive than the older approaches, so it will be my preferred approach for such work – for the time being, at least.

Image Credits

"La Gioconda" (aka "Mona Lisa") by Leonardo da Vinci *http://en.wikipedia.org/wiki/File:Mona_Lisa,_by_Leonardo_da_Vinci,_from_C2RMF_retouched.jpg*

***Mark** uses Inkscape to create comics for the web (http://www.peppertop.com/) as well as for print. You can follow him on Twitter for more comic and Inkscape content: @PeppertopComics*

RETURN TO CONTENTS

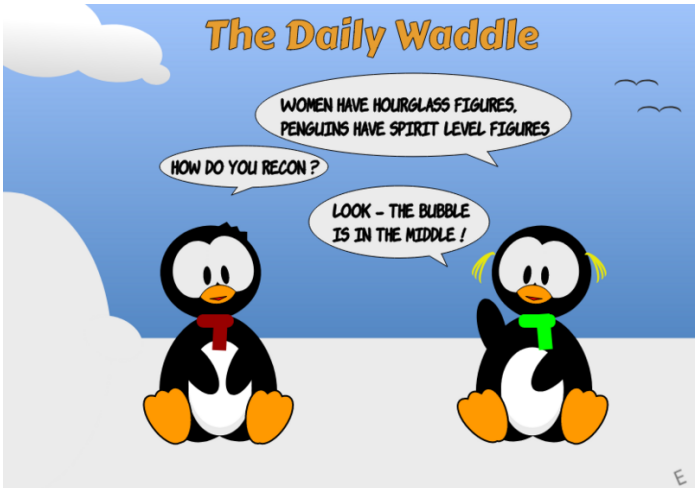
The Daily Waddle
WOMEN HAVE HOURGLASS FIGURES, PENGUINS HAVE SPIRIT LEVEL FIGURES
HOW DO YOU RECON ?
LOOK - THE BUBBLE IS IN THE MIDDLE !

RETURN TO CONTENTS

shells.com
Linux on Your iPad
For as low as $4.95, you can have your own personal Linux cloud computer in minutes on any device.
GET IT ON macOS
Get it on Windows 10
Download for Linux
Download on the App Store
GET IT ON Google Play

RETURN TO CONTENTS

BODHI CORNER

Written by Moss Bliss

Welcome to the first entry in what I hope will be a continuing series, Bodhi Corner. I've been using Linux since 2002, with full time use from 2009-2012 and again from 2015 to the present. I will have a number of insights into Linux and, specifically, Bodhi Linux, not all of which will be 100% "real". I'm an older autistic individual, and sometimes my perceptions turn out to be a bit sideways, but I'm sure I will get more factual and informative as the series progresses.

I was becoming a fan of small Linux distros for some time, being all in on Puppy from the time before Lucid Lynx. At one point, someone came out with MacPup, which was Puppy but

with the Enlightenment E17 desktop. I had two versions of MacPup before the distro stopped publishing, and was hooked on E17 as my favorite desktop.

I went looking for other distros featuring e17 (technically, enlightenment 0.17), and had trouble finding any. I ran into a disk for Bodhi 3.1 in Linux Magazine, but could not manage to make it install on the equipment I had at the time. I kept track of it, as nothing better was coming along, and when Bodhi 5 came out (the last edition featuring Jeff Hoogland as the lead dev), it finally worked on the machine I had. I have used it ever since.

Bodhi 5.1 was a huge step forward, as it was the first version where Robert Wiley served as lead dev, as he continues to do to the present day. It is also when the fork of e17, Moksha, which served as the desktop or window manager for Bodhi, started growing beyond e17 but away from the main ‘e’ development.

Why Moksha? Because I am addicted to Choice. In desktops, you have your choice of either Gnome or other GTK variants (Xfce, MATE, Zorin) or KDE Plasma or other Qt variants (Trinity, LXQt). The originals of both COSMIC from System 76 and Budgie from Solus were just GTK rearrangements. Until System 76 started writing their new version of COSMIC in Rust, those were your choices...

Unless you were aware of EFL (Enlightenment Foundation Libraries) and the enlightenment window manager.

Enlightenment was dubbed, early and often, the "original eye-candy window manager". It wasn't just a rehash of what you'd like your Windows PC to look like, and it wasn't, really, a full desktop. But it looked and worked like one. Tired of moving your mouse around to find the menu button? In ‘e’, you can just click anywhere on the desktop and get a menu.

So, over the years, I got more and more interested in Bodhi and the Moksha desktop. The current version of Moksha is still based on e17, but has gone far beyond that, both in stability and features. And in early 2025, two features of ‘e’ which had been forgotten have been added back to Moksha, but I'll talk about those later.

In 2021, I received some money from the death of my mother, and decided to spend some of it donating to Bodhi. I was already considered a team member, but had not really done anything with that. I have continued making regular monthly donations, and even supplied Robert Wile with a Lenovo ThinkPad T540p to continue developing on (his old computer, from 2008, was on its last legs). Just last month, I started adding Bodhi content to mintCast, where I have been a team member since 2018, and now I have been given the opportunity to talk about it here in Full Circle Magazine.

Basic information. Bodhi is based on the LTS versions of Ubuntu, but with Moksha Window Manager in place of other desktops. The installer is Ubiquity, the package manager is APT, and Synaptic has always been available. The file manager seems to toggle from version to version from PCFileManFM to Thunar. Bodhi is packaged in 64- and 32-bit ("Legacy") versions. The most current "official" 32-bit is still version 5.1, but there are fully functioning betas of both Bodhi 6 and Bodhi 7 (based on Debian Bullseye) for 32-bit. The 64-bit version comes with your choice of the LTS kernel, the HWE kernel, or even the System 76 kernel (newer than HWE, better for gaming on new equipment). There is also a Beta3 version of 64-bit Bodhi 7 based on Debian Bookworm, which works quite nicely and is often referred to as DeBodhi.

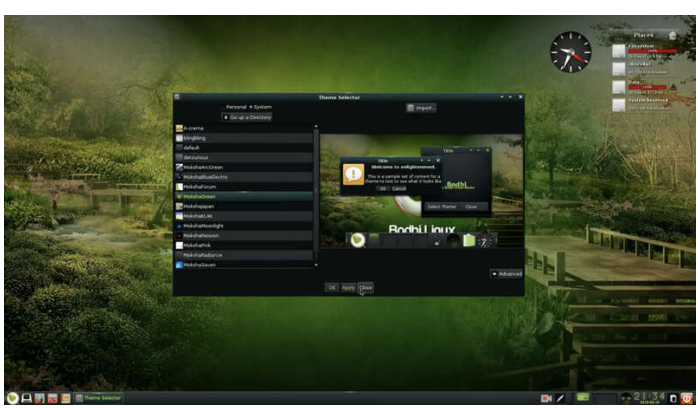

The software included is fairly minimal, although they do

produce a version called AppPack which is pretty full-featured. The basic version of Bodhi comes with the following software choices (although you can add anything Debian- or Ubuntu-related as you like):

- *Terminology Terminal Emulator*
- *Chromium Web Browser*
- *Thunar File Manager with archive plugin*
- *Leafpad Text Editor*
- *ePhoto Image Viewer*
- *aRandr Monitor Settings*
- *Web Browser Manager*
- *Engrampa File Archiver*
- *Pavucontrol Pulse Audio Control*
- *Gnome Language Selector*
- *Synaptic Package Manager*

You can go to *https://www.bodhilinux.com/w/selecting-the-correct-iso-image/#Legacy_32-bit_only* and check out what applications are in the AppPack version.

You can see some of the changes just on the screenshots provided from different versions.

The latest updates to Bodhi 7 (and the under-development Bodhi 8) include the new wallpaper picker, called Wallscape, and the return of the Moon module (shown replacing the Bodhi logo on my own desktop screenshot). Work is in progress on the Drawer module, which I will talk about more in future articles as it develops.

I should wrap this up for now, but will be back with more in subsequent issues. I look forward to presenting more - and more technical - information on Bodhi next month.

RETURN TO CONTENTS

UBports Devices: News From The Community

Written by UBports Team

Focal 20.04 OTA-8 Released

We are very pleased to announce the roll out of Ubuntu Touch OTA-8 Focal. The eighth stable release based on Ubuntu 20.04 brings a few improvements and changes. All the details can be found in the OTA-8 blog at:

https://ubports.com/blog/ubports-news-1/post/ubuntu-touch-ota-8-focal-release-3953

RETURN TO CONTENTS

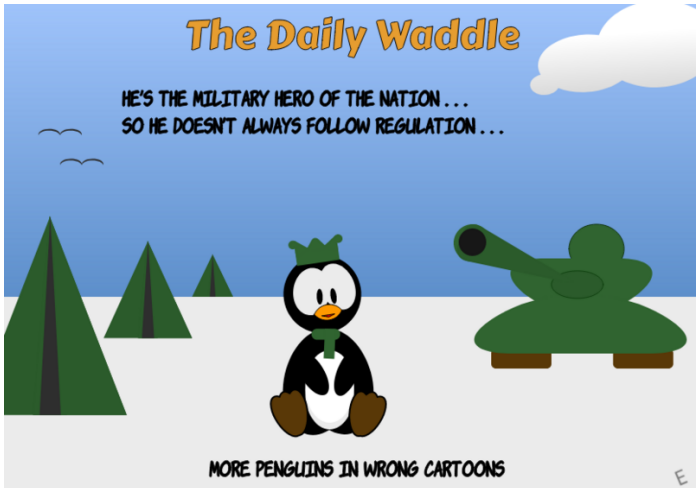
The Daily Waddle
HE'S THE MILITARY HERO OF THE NATION . . .
SO HE DOESN'T ALWAYS FOLLOW REGULATION . . .
MORE PENGUINS IN WRONG CARTOONS

RETURN TO CONTENTS

Your trusted source for fast learning

in easy steps

Clear and concise guides to help you learn with minimal time and effort!

Order direct from ineasysteps.com **and get 25% discount!**

Enter **FC25** in the Coupon Code box at the checkout.

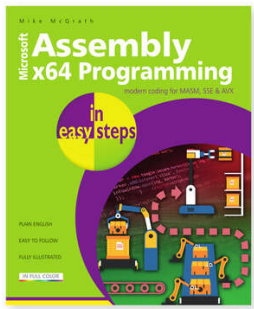

Assembly x64 Programming in easy steps teaches how to write code to create your own computer programs in the modern x64 Assembly language on Windows. Written with beginners in mind, it shows how to:

- Store and manipulate data in 64-bit registers.
- Control program flow.
- Create reusable blocks of code in program functions.

This primer includes demonstrations of parallel processing with 128-bit Streaming SIMD Extensions (SSE) and 256-bit Advanced Vector Extensions (AVX).

Complete example programs demonstrate each aspect of coding, together with screenshots that illustrate the actual output when each program is executed. The code covered within the book is color-coded, making it even easier to grasp.

Includes free, downloadable source code to get you started straight away!

£11.99 / $15.99 paperback (ebook version also available)

ISBN: 9781840789522 / 192 pages / By: Mike McGrath

RETURN TO CONTENTS

My Opinion: To Rust or not to Rust

Written by Alan Ward

This article has been inspired by the recent controversy about the inclusion – or not – of some code written in the Rust programming language inside the Linux kernel.

For some context, I should mention that the Linux kernel, like most other UNIX-like kernels, has been written mostly in a standard dialect of C. To sum up, the basics of the question is that some developers are in favor of including code in Rust in, at the very least, parts of the kernel source code. Device drivers have been mentioned. But there has been some pushback by other developers who do not agree with this move, in at least one case leading to a discussion about whether code written in Rust should even be able to access routines written in C. Linus has come in very clearly on the topic, setting out that developers responsible for parts of the kernel code who wish to write in Rust may do so, while others may prefer to stick to C. But people must be able to access each other's code with no hindrance.

There are arguments to be made on both sides of this matter. On the one hand, adding an additional programming language to an already complex kernel source code tree further increases complexity. There are no two ways about that.

Also, being able to read and understand the complete source code will be made more complex, since now a working knowledge of both languages would be required. However, it must be said that Rust syntax is actually rather close to C, and, if anything, is slightly more readable. So this argument is actually not really valid, in my view. On the other hand, Rust is presented as a much more secure programming language than C, specifically as regards memory management. It could be said

that Rust has been designed to avoid the typical programming errors that plague beginners coding in C, such as assigning pointers but then not freeing the memory they point to afterwards, or inversely trying to free up too much space, or fringe effects such as modifying what a pointer points to in two code sections at the same time. This is said to be literally impossible in Rust, or, at the very least, much harder to make mistakes. The Rust compiler has very many built-in safeguards to avoid such potentially bad and hard-to-detect mishaps in a critical piece of software as is the kernel.

From the point of view of us *buntu users, what exactly could the potential effects of including more Rust code in the Linux kernel be? These are not very many, in reality. Perhaps the kernel itself would be slightly larger, since it is unavoidable that some routines should be duplicated. We can expect some degree of increase in stability, though it must be said that the current state of affairs is not bad at all from that point of view. What will certainly be impacted is the toolkit one would need to compile a kernel. In addition to the standard GCC compiler suite, one would now need the Rust toolchain as well. In some cases, this may be something of a drawback, such as when porting the kernel for a new physical platform. Preparatory steps will now include not only making sure some version of GCC is available to cross-compile for the new architecture, but we may very well find ourselves also porting the Rust compiler first, or at least some parts of it. Since the Rust compiler uses LLVM for code generation (see the “Rust Compiler Developer Guide”, available online), it may very well be that the complete Linux kernel project may find it best to slowly go from the GCC compiler to LLVM. There already seem to be projects such as Android or ChromeOS that use this latter toolchain, so it could very well be not only possible, but even relatively easy. Once more, this would probably be an issue that most users would hardly even be aware of.

So, in final analysis, is all this merely but a storm in a teacup? Not really, there is more to it.

For those who are not programmers, I would like to draw a

parallel between the different philosophies of the C and Rust programming languages, and using a manual or an automatic transmission in a motor vehicle. People who are used to shifting their own gears by hand are usually rather good at it and may even take some pride in the fact. After some practice, they may gain finer control over their car's drivetrain and a better feeling for the forces moving through the shafts and into the wheels, and, in general, a better comprehension of the mechanics of shifting gears. This would in general not be the case of the driver of a car who uses an automatic gear shift, and who perhaps is prepared to sacrifice some mastery of the mechanics in exchange for the convenience of just getting from one place to another with minimum fuss. In much the same way, a proficient C programmer will have gained the experience, over the years, to not make specific mistakes. For this person, the Rust compiler's very many nanny-helping mechanisms will not be much of a practical help, and may actually constitute a hindrance for getting things done in an efficient manner. But a Rust compiler will avoid many mistakes and may actually be easier on younger programmers, empowering them to build complex pieces of software without having to previously take some hits over arcane types of bugs.

Is this a good thing? Perhaps so, and perhaps not. But consider: we are speaking of open-source development that is mostly done on a voluntary basis. The average age of contributors may be difficult to ascertain, but some key figures of this world are clearly not going any younger. When current leaders need to step down and hand the reins to the next generation, we will probably have to talk about programming style and whether the more hard-core C language with pointers and whatnot is really suited to the newer programmers' style or vision of kernel programming. In that sense, Rust may very well be a suitable middle ground to build a bridge between generations. It still retains some of the quirks of C, but there is now a handrail between the walkway and the drop down into the abyss of kernel panic.

Also, have you noticed how some youngsters nowadays are quite incapable of handling a manual transmission properly?

***Alan** holds a PhD in Information and the Knowledge Society. He teaches engineering at Escola Andorrana de Batxillerat (high-school). He teaches Operating Systems at the University of Andorra and has previously taught GNU/Linux systems administration at the Open University of Catalonia (UOC).*

RETURN TO CONTENTS

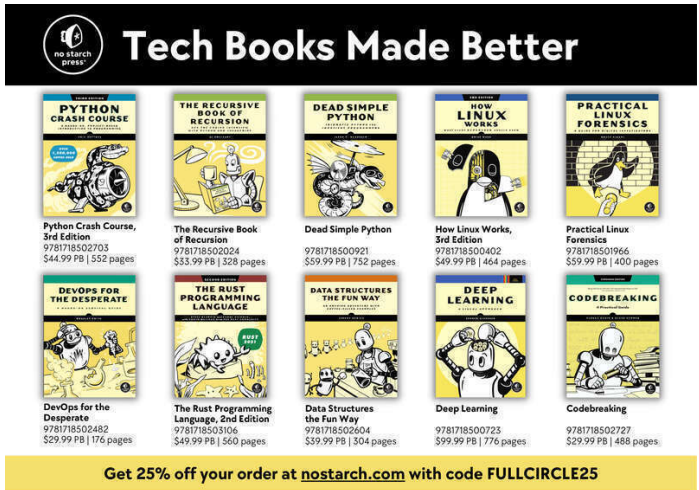

RETURN TO CONTENTS

HowTo: Write For Full Circle Magazine

Written by Ronnie Tucker

Guidelines

The single rule for an article is that it must somehow be linked to Ubuntu or one of its many derivatives—Kubuntu, Xubuntu, Lubuntu, etc.

The Official Full Circle Style Guide can be read at: *http://bit.ly/fcmwriting*

Please read this document before submitting an article. Follow the guidelines and you will have a much better chance of seeing your article in Full Circle.

Writing

There is no word limit for articles, but be advised that long articles may be split across several issues. In your article, please indicate where you would like a particular image to be. Please do not use any formatting in your document.

Images

Images should be no wider than 800 pixels, in JPG format, and use low compression. When you are ready to submit your article, please email it to: *articles@fullcirclemagazine.org*

Non-English Writers

If your native language is not English, don't worry. Write your article, and one of the proofreaders will read it for you and correct any grammatical or spelling errors. Not only are you helping the magazine and the community, but we'll help you

with your English!

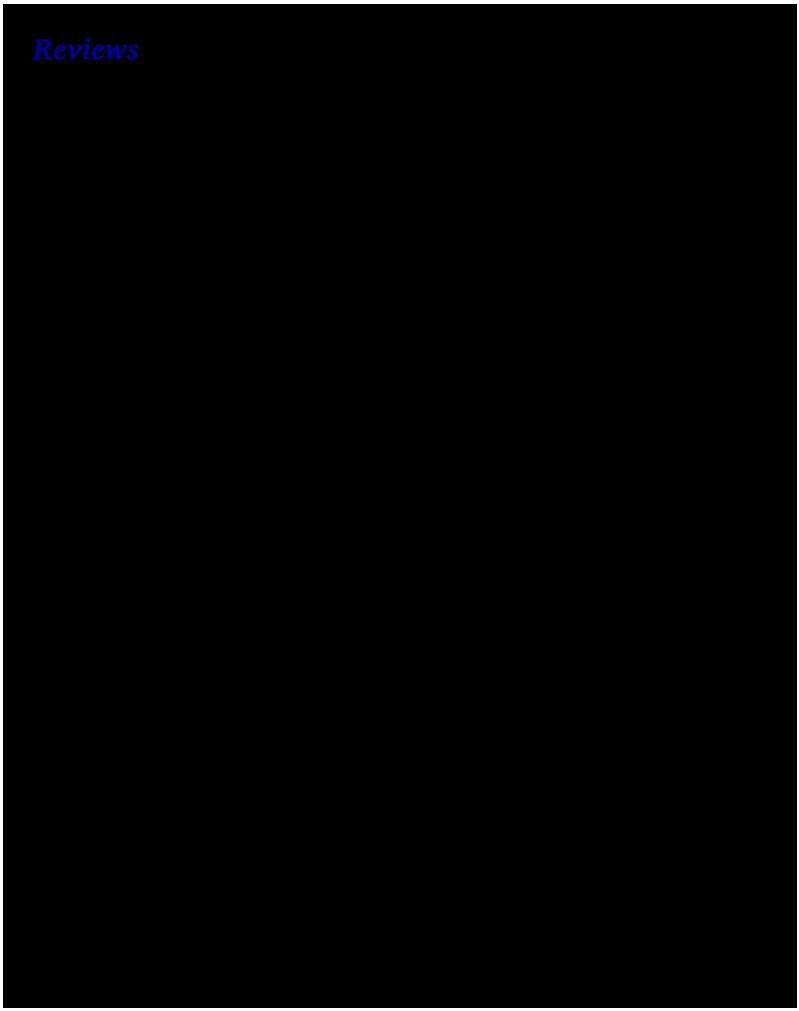

You don't need to be an expert to write an article – write about the games, applications and hardware that you use every day.

RETURN TO CONTENTS

Review: DistroSea

Written by Art

I admit it, I'm a distro-hopper. I doubt it's a stretch to estimate I've tried at least 100 Linux variants over the past 20 years. Normally, that means downloading an iso image, creating a bootable USB, firing up the virtual machine app (VMA), and going from there.

Great if your computer has the hard drive space and horsepower to handle such work, not so hot otherwise. It's also time consuming with roughly an hour to complete the whole process (plus another hour or so figuring out what you tested was either garbage or gold).

For a while (around 2020, but nobody seems quite sure), there was a website called DistroTest.net that allowed you to try various operating systems via an online virtual machine (Qemu), but it disappeared roughly 3 years ago and the website is now occupied by a foreign gambling venture (appropriate given that either site was a gamble).

DistroTest offered 800+ Linux variations (some of them repeated with different desktop designs) and, as I recall, others were outdated or nonfunctional when I tested them (hence the motto, "Test it before you hate it"). Might see names you'd recognize, yet others remained in the distant shadows of the Linux world.

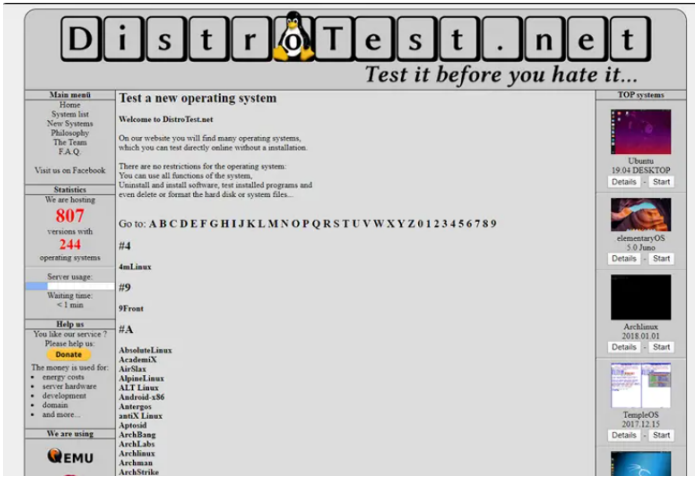

It did have a certain charm about it, though. Using Qemu as its testing platform, you could, OS permitting, take a peek at what was being offered in the Linux community. Bizarre and unknowns were treated equally there.

After they disappeared (and I do mean disappear - no adios, au revoir, nothing), I just went back to running iso images off USB drives and thought nothing more of it.

I recently heard about DistroSea.com, a similar testing site. I don't believe they are the same company but I can state the website is better organized, at least visually (*https://distrosea.com/*). Instead of here, there and everywhere listings, you see icons for the following (all are Linux, so I'll dispense with adding that word unless it's part of the name):

Alma	DragonflyBSD	Lubuntu	Devuan
Slackware	Alpine	Edubuntu	ageia
Slax	Antix	ElementaryOS	anjaro
SliTaz	Archcraft	EndlessOS	X Linux
Solus	Arch	Fedora	Nitrux
SparkyLinux	Arco	FreeBSD	NixOS
Artix	Garuda	OpenIndiana	Tiny Core
BlendOS	Gentoo	OpenSUSE	Triquel
Bodhi	Haiku	PeppermintOS	VanillaOS
BunsenLabs	HoloISO	Pop! OS	Void
CachyOS	KDE Neon	Porteus	Xubuntu

CentOS	KolibriOS	PrimTux**	ZorinOS
Chimera	Kubuntu	PureOS	Debian
Linux Lite	RebornOS	Deepin DE*	Linux Mint
Rocky	Linux Mint	Siduction	
	Debian		

Notes:

- ** Chinese only, I let the website administrator know since there is an English version available.*
- *** French only, but the description clearly states that. For the picky readers out there, I know that some of these OS titles are not capitalized, but I'm listing them as the DistroSea website does.*

Ubuntu was left out because it gets 6 additional entries for Budgie, Cinnamon, Kylin (Chinese only), Mate Studio and Unity.

As you can see, nearly every type of Linux OS is covered - Arch, Debian, Ubuntu, Gentoo, Slackware, BSD, and probably a couple I missed. Desktops represented include Pantheon, LXQt, Enlightenment, Deepin DE, Budgie, LXDE, Mate, Gnome, XFCE, Unity, and so on.

Maybe not as well stocked as DistroTest was, but at least these, so far, have turned out to be active OS offerings.

And, from what I can read on their Facebook page, this listing is ever changing so you can expect to see some obscure or long forgotten titles popping up.

For each of these distros you'll find the appropriate icon as shown below (which is a small screenshot of a much larger page).

To use DistroSea you have the option of signing in, but those who do not cannot use the browsers or app stores while testing - that is reserved for those who use their Google account to sign in (I tried both ways and saw no difference).

Users have their choice of language to include English, Spanish, German, French, Italian, and Portuguese; however, any choice you make influences the DistroSea website, not the OS. Choose Spanish as your language of choice and the description will show in Spanish but the OS itself will show and run in English (although users can change the language in many live versions, but not all).

That applies to any language chosen with the exception of PrimTux (the OS is French - period), Ubuntu Kylin (Chinese), and DeepinOS (Chinese).

Beyond that language mystery, clicking on any of the icons results in the following:

1. *A brief description of the OS along with desktop options, if any. You may be asked to choose from Gnome, Mate, Cinnamon, etc. In addition, you may find several versions ranging from old to new, stable and unstable.*
2. *You'll then be asked to click on a box to prove you're human. Fortunately, you won't have to pick out bridges, fire hydrants, or whatever.*
3. *Then comes the dreaded queue. No matter how many users it showed ahead of me, I never waited more than a minute with the longest waits being for Nitrux (no idea why) and nearly*

any form of Ubuntu.
4. *When your turn is up, you'll then be asked to click on Continue and off you go.*
5. *When you're done, power off the OS and you'll be returned to the start page where you can choose another OS or leave.*

Since you're using DistroSea resources, you won't be asked for RAM assignment, login credentials (unless the OS requests such and DistroSea will let you know), or how much drive space you want to sacrifice. These are all live versions.

As for OS speed, I'll just say any you pick will act just like you'd expect a live version to - slower than normal. I asked the developer via email and Facebook about what VM app was used (even though I suspect Qemu) and RAM allocated, but got no response either way (the Facebook page hasn't been updated since October, 2024).

Unlike VirtualBox or Gnome Boxes where you can allocate RAM and drive space to any tested OS, you're at the mercy of resources provided by DistroSea.

For the most part, I found DistroSea to be a worthy alternative to downloading multiple iso images, especially if you offered desktop options and don't have a clue what the difference is between Mate and Cinnamon or Gnome. And, since you're using their resources, you don't have to worry about sacrificing RAM or drive space to a VM.

Where this comes in handy is when an OS offers multiple desktop choices, but not in the same package. SpiralLinux offers Mate, Gnome, XFCE, Cinnamon, etc, for choices, but you'd have to download each iso image to test all of them. With DistroSea, pick your package choice, play with it, and then move on to the next one.

On that same topic, DistroSea saves the hassle of having to wait for downloads. Using SourceForge as its download source means getting multiple copies of the SpiralLinux OS can be a painfully long affair (nearly a half hour for me - and I have a high speed connection).

It has also let me discover that some distros are smoke and mirrors and really stink during testing. Yet others I tried, like RebornOS and Siduction turned out to be better than I anticipated.

However, nothing is perfect and DistroSea does occasionally bog down and I can only surmise it's because users are throttling the resources available. And, in some cases, I think the OS I was trying just failed. Ubuntu Cinnamon was a big flop - twice. Apps failed to open and I got failure warnings galore.

Then you have the uncooperative OS that refuses to shut down. That's where the little pop-out below comes into play.

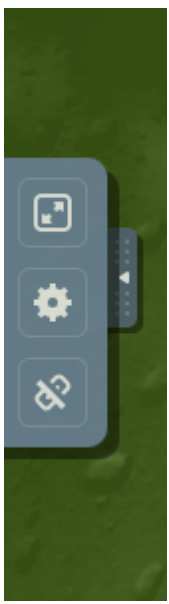

The top part is for expanding the main window, the settings gear is there to tweak the incidentals, and the bottom is to kill the signal when the OS just won't take a hint and quit when you tell it to.

And what happens if you don't shut down or disconnect? You'll get a reminder that the OS is still running and that session will be cancelled if you attempt to pick another OS to run.

Why did I include that little sidebar screenshot? It's nearly identical to the one DistroTest used. Main reason: I'm sure DistroSea is using Qemu, just like DistroTest did.

Remember earlier when I made a comment about app stores? Of all the live versions I tried (at least 10), not one would allow me to install apps and all but one locked up when trying to review apps; however, when I ran most of those via my USB stick as a live boot, I could use individual app stores (although, of course, everything disappears upon shutting down unless you're using a persistent drive).

There are small goofs here and there. For example, DeepinOS from China is downloadable in English from their website, but the DistroSea version is in Chinese with a smattering of English here and there (good thing there is an international symbol for shut down).

However, DistroSea is worth a try now and then whenever I see a new OS, and it allows me to do a brief look around before deciding to download the iso image for additional testing.

RETURN TO CONTENTS

Review: VirtualBox vs Gnome Boxes

Written by Art

When it comes time for testing Linux OS variants, my go-to virtualization machine software (VMS) for over a decade has been Oracle's (formerly Sun Microsystems) VirtualBox.

I rarely encountered any major problems; however, in the past few months, I've had kernel compilation errors where none exist, and have discovered file fragments from operating systems long deleted. Part of the issue may be an update issued by Oracle in October, 2024.

Online forums were no help because Oracle claims once you delete the secondary OS, or VM (virtual machine), it should leave no fragments which, unfortunately, and obviously, it does.

Time to find a new VMS.

In researching I found alternatives:

- *Linux KVM (Kernel Virtual Manager), aka Virtual (or Virt) Manager, aka QEMU/KVM. Once components are added properly, (see https://ipv6.rs/tutorial/Linux_Mint_Latest/KVM/ for info on that), you can run virtual machines as if they were actually installed on your pc as the main OS (which allows them to run faster than other VMs).*

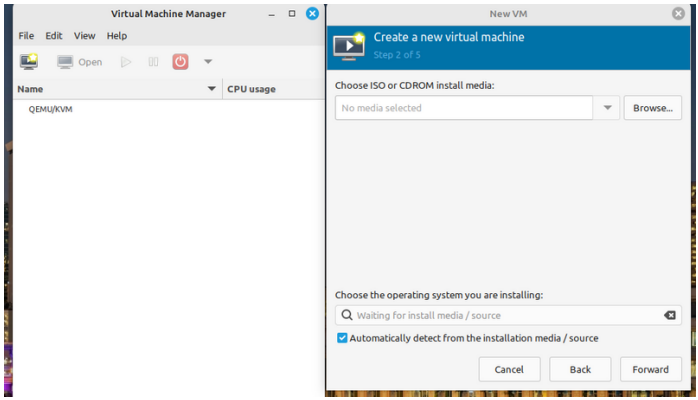

However, in use it's a bit more complicated than some may desire (more on that later); however, there is no denying it does the job admirably.

- *Broadcom's VM WorkStation Pro, not to be confused with VMWare Fusion Pro (the former is for Windows and Linux, the latter for Macs). Originally a proprietary VMS, it went free for personal use earlier this year, or so Broadcom states.*

However, the website (vmware.com) makes no mention of the product (they prioritize cloud solutions now), and other sites redirect a sign-in for a Broadcom account. Don't have one? Go ahead, sign up - I dare ya. You'll succeed in getting plenty of junk emails and a "404 Page Not Found" notification every time you attempt a download.

That's why you won't see a screenshot here - I couldn't get it to download!

The worst part is… at one time it was a decent, if proprietary, VMS.

- *Gnome-Boxes, often seen as Boxes in many software managers. Of these three, it has an easy-to-use interface and is nearly impossible to goof up.*

But there is something of a weird past for Boxes. Released in 2011 by Red Hat in conjunction with Gnome, the last copyright on it is dated 2018. In 2019 IBM took over Red Hat and Boxes

has been maintained for Gnome by Felipe Borges of Czechia with the last update being issued in September, 2024.

That results in a bit of confusion seeing that the most recent version of Boxes from 2024 carries a copyright date that ends in 2018 only because Red Hat became an IBM subsidiary (see screenshot below, dated December, 2024). To the unaware, it looks like the app is outdated by 6 or so years, but it's actually up-to-date, it just doesn't say so.

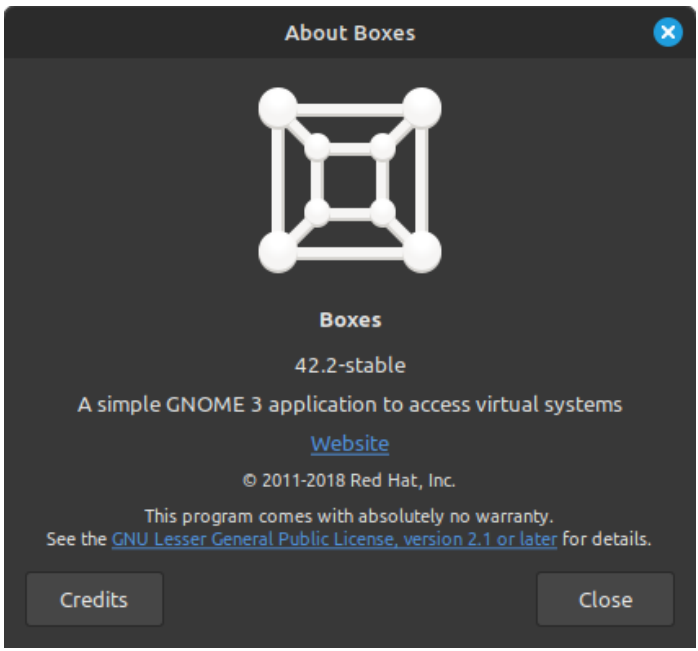
About Boxes
Boxes
42.2-stable
A simple GNOME 3 application to access virtual systems
Website
© 2011-2018 Red Hat, Inc.
This program comes with absolutely no warranty.
See the GNU Lesser General Public License, version 2.1 or later for details.
Credits
Close

Of the three, it's clear Boxes and KVM are the only two worthy of further consideration, so it's off "a testing" I go.

Installation

Boxes: While all major software centers and package managers carry Boxes, you're better off using sudo apt install gnome-boxes in the terminal. Seems that some other versions may be outdated or missing components.

And what about going to the Gnome website (*https://*

apps.gnome.org/Boxes/) and getting a copy there? It's a buggy Flatpak version that throws kernel errors. Avoid it!

Download is roughly 60 MB and expands to 360 MB once installed.

Virtual Machine Manager:: Installation is just as uncomplicated if you opt to go through your software manager and use the Flatpak version; however, I've seen some reports stating that that version is problematic.

It might be better to follow the link I posted above (*https://ipv6.rs/tutorial/Linux_Mint_Latest/KVM/*) and follow the outlined terminal command and installation method.

If you follow the procedure, you'll have a new menu entry titled Virtual Machine Manager (don't bother looking for QEMU, KVM or any combination thereof - it won't show).

Download size is about 25 MB and expands to approximately 75 MB, a substantially smaller footprint than Boxes. This is due to it being treated as part of your main Linux kernel system instead of an addition to it. In a sense, it's not an app, it's an extension.

Finding Your First VM

Boxes: It gives users the choice of using their own ISO image or they offer to download an image from a built-in catalog of Linux variations.

By default, Boxes highlights Ubuntu versions 20.04 and 24.04 (what they call "Express Installation") plus there's a dropdown box that offers several other Linux variations.

Click if you dare, because the use of "express" is greatly exaggerated. What might take a minute or two on a high speed connection can take 30 minutes or more using what they offer.

Believe me, just download your own iso image and save the frustration.

Virtual Machine Manager:: It requires you to bring your own image to the game.

Both: And where should you park your image once you have it? While Virtual Machine Manager prefers you put it in their designated file, it doesn't matter. When it comes time for installation, you can point it in the general direction, even to an SD card, and it'll find it.

Boxes doesn't care. Park the ISO image in files, on your desktop, or an SD card. Mouse-click the file and it'll find it.

ISO Installation

Now comes a real test - installation of an ISO image which I assume you've downloaded.

Boxes: It requires you click on the upper left X at which point you'll see this:

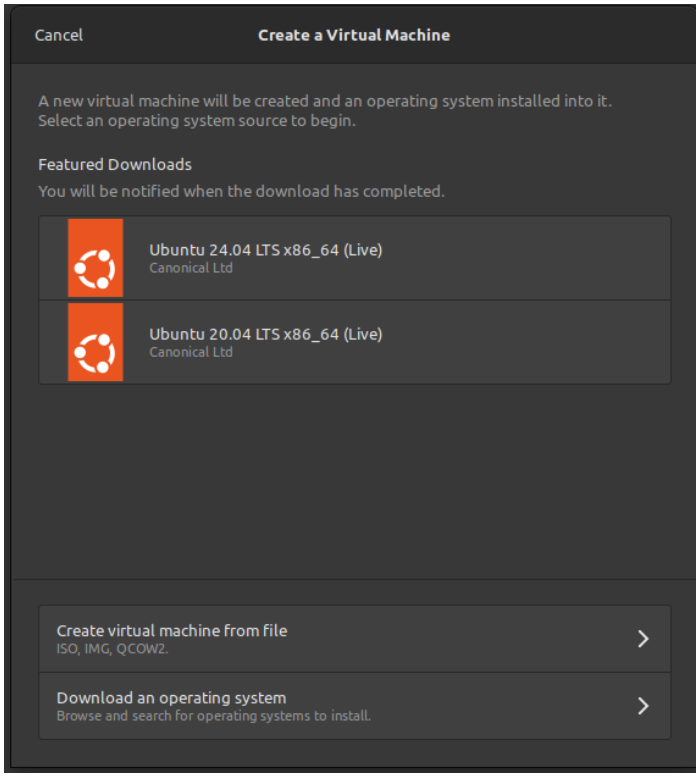
Cancel
Create a Virtual Machine
A new virtual machine will be created and an operating system installed into it.
Select an operating system source to begin.
Featured Downloads
You will be notified when the download has completed.
Ubuntu 24.04 LTS x86_64 (Live)
Canonical Ltd
Ubuntu 20.04 LTS x86_64 (Live)
Canonical Ltd
Create virtual machine from file
ISO, IMG, QCOW2.
Download an operating system
Browse and search for operating systems to install.

There must be an Ubuntu connection somewhere, because those are the first two offerings you'll see, but you'll want to click on the third choice, "Create a virtual machine from file".

After finding the ISO, Boxes sets RAM at 2 GB minimum and HDD space up to 25 GB; however, both figures are adjustable.

But there's a kicker, Boxes has a tendency to misread an OS title at which point it assumes it's Ubuntu 24.04 or 20.04 (definite Ubuntu bias); however, click on either one to start the loading process and it'll figure out what the actual OS is. Goofy but true.

I tried Edubuntu in Boxes and it didn't initially recognize it as an Ubuntu derivative. Really? Did it not see "ubuntu" after Ed? I merely clicked on Ubuntu 24.04 and within a few seconds it figured everything out.

From there on, the ISO will load and you should have a running test OS.

Virtual Machine Manager:: I wish it was the same as Boxes, but its process is a bit trickier and a lot messier, to wit:

1. *Click File and New.*
2. *Browse for ISO image.*
3. *Designate Linux by unclicking, "Automatically detect…." and typing Linux in that area. For some reason it can't detect Linux iso images, even if Linux is in the title.*
4. *Chances are, it will still misread everything at which point you'll have to use the dropdown box and choose Generic Linux 2022 (clear down at the bottom of the listing). Believe me, make life simple and go this way. Should you choose the wrong Linux title in between it could get messy.*
5. *Pick RAM and number of CPUs to use. Unsure? Just pick what it offers.*
6. *Create Virtual Hard Drive (it starts at 25 GB). Using a live version and don't care about drive space? Pick 25 GB anyway. It won't do anything unless you physically install the OS as a VM.*

If all goes well, your ISO will load but you just wasted several minutes going through the routines above.

Fortunately, once you've done it, future attempts will be less cumbersome.

Allocation of Resources

Just want to peek under the hood of a new OS without having to install it and wasting limited drive space?

Don't install the OS - create a VM using a live version.

Live systems operate out of RAM and, therefore, disappear once the session is over; however, you must still download an iso image and that'll probably be around 2 to 6 GB in many cases.

If you're wondering, most Linux offerings I've tested use live

versions.

The only disadvantage here is that they vanish after every session and nothing is saved of your progress. If you're just looking around, it makes no difference anyway.

However, if you desire to install an OS, keep in mind it will take a sizable chunk out of your hard drive.

If you're seriously strapped for drive space but still want to do VM testing, consider parking your test OS on an external drive. Both Virtual Machine Manager and Boxes can run off external drives.

Does that help? Yes, it does save your drive space by putting VMs on an external device; however, it also rates a negative in that running VMs off an external drive through a VMS is, as you might say, sluggish.

Inadequate RAM can be problematic . If your PC is pushing 1 to 4 GB RAM and that's the best you can muster, avoid doing anything with VMs. Just not enough horsepower. The VM might work, but it'll stress both you and your computer.

Why? When running a VM, it must borrow RAM for that secondary OS to run properly. Take too much and it could cause your main OS to lock up (I've done that). Don't have enough and neither OS will run properly.

Interface

Boxes: It's from the school of minimalistic aesthetics. It's black-and-white, with a dot of blue in the upper right corner.

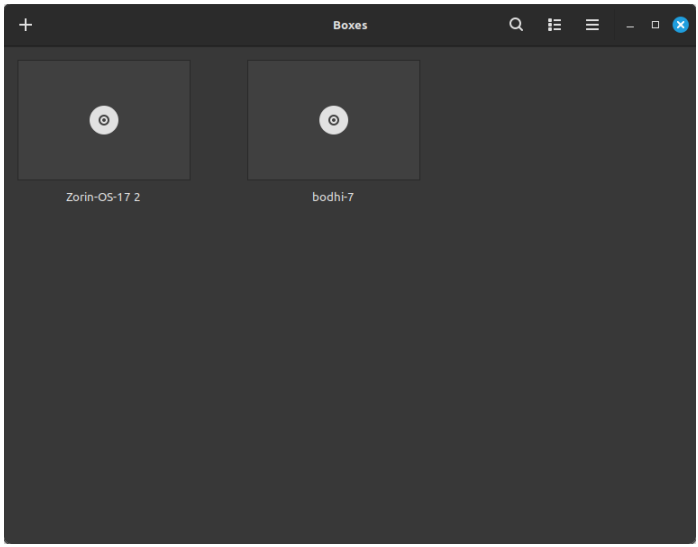

Boring, but functional.

If you look at the example above, you'll see two boxes occupied by my test operating systems (Zorin and bodhi). If those weren't there you'd just see an empty black box with an instruction banner to click on the + to add a new box and OS. Another icon is for search, while the lines to the right of that changes your VMs from boxes to a standard title listing. The triple lines bring up a menu for Keyboard Shortcuts, Help, and About Boxes.

Worth noting, users can opt for a vertical listing if they don't like the look of squares or boxes.

Virtual Machine Manager:: It's stark but functional with just enough icons and headings so that the casual user can figure it out.

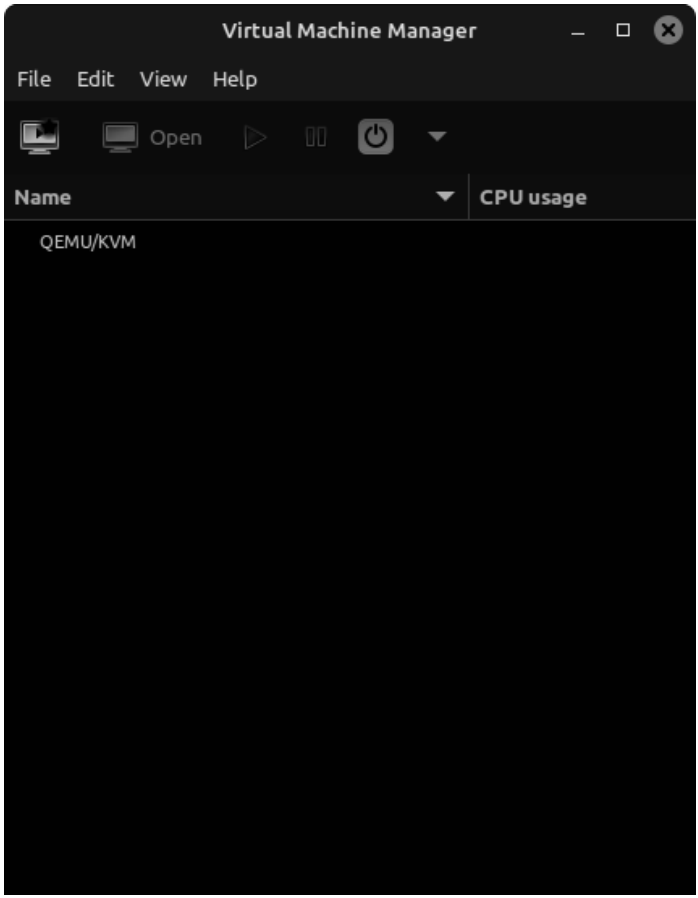

So stark, I had to change the color to black because it's mainly white-with-a-dash-of-color and the white part disappeared into the background. Look at my first image for a true representation.

Otherwise, the menu is straightforward without much drama, and the included CPU meter is handy for telling you when an OS is getting a bit too ornery.

Need Instructions?

Boxes: It has an online manual at *https://help.gnome.org/users/gnome-boxes/3.26/*, and is nicely organized and not quite as complex as VirtualBox (which has a nearly 400 page online

help document that is longer than most textbooks on Linux I've read). Still a lot to digest, but easy to navigate and most of the instructional chapters don't seem to occupy more than a couple pages.

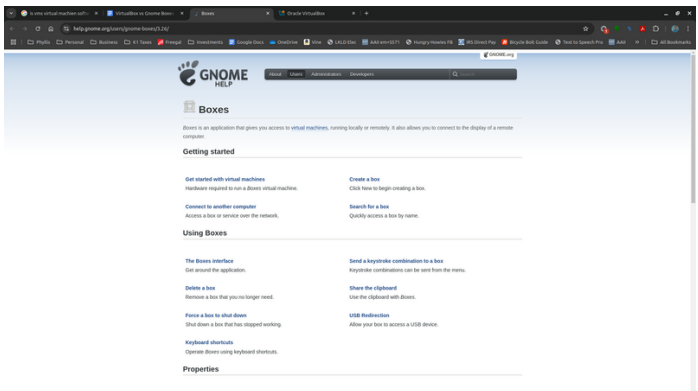

Virtual Machine Manager:: You can go to *https://virt-manager.org/* whereby you'll be greeted by, well, nothing much.

Or you can go to *https://ubuntu.com/server/docs/virtual-machine-manager* where you'll see, well, not much more.

I find online forums to be especially helpful.

VMS Issues

Whether you use Virtual Machine Manager or Boxes, you can expect the following:

- *Decreased speeds. Any OS running in a VM will be slower than it is in realtime on a physical computer. You're running off borrowed resources.*
- *Want to check out the software center on your chosen OS? Have fun. Might not even open or give you the time of day.*
- *You'll probably notice a lack of a wireless (or wired) signal icon when running your VM. The VM is piggybacking off your computer's resources - not using its own.*

And, in some cases, it may show you using a wired connection when you might be using wireless.

For example, Deepin showed I was running off a wired connection when, in fact, I was running a wireless connection.

And SparkyLinux didn't even show a wireless or wired icon at all, but yet I could use its browser while in VM mode.

It's not where you think it is. Just because you have pictures on your computer doesn't mean your VM OS will be able to find them. In fact, I can pretty much guarantee it.

There is a workaround to this, but it often requires you install the VM and not use live mode. In this scenario, you can possibly share files between the VM and your main OS.

Somewhere down the line, things will get wonky.

Both Boxes and Virtual Machine Manager showed a nasty tendency to bog down the longer I used any OS. In one case, Deepin attempted to do an update and that spiked CPU usage so high it locked up the VMS, Deepin and even killed my computer's WiFi connection.

And don't be surprised to hear your computer fan running at double speed. Lots of CPU stress in some cases.

Shutting Down/Ending the Session

Boxes: It requires you shut down the OS in the VM at which point the box goes blank, like you saw in the previous screenshot. What if you decide to shortcut and just click the X to terminate Boxes and not the VM first? Bad human! Now the VM is running in the background sucking up your RAM and CPU resources. Instead of a blank, black, box, you get one that's animated. In the screenshot below, although it's not obvious, that OS (SparkyLinux) is running in the background even though I don't have the VM open.

Shut down the running VM OS first, then shut down Boxes.

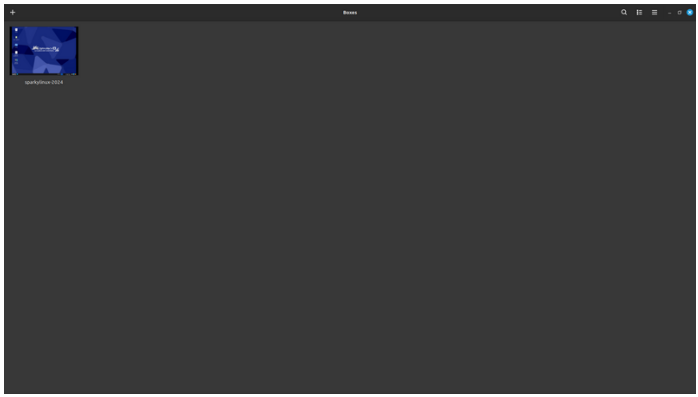

Virtual Machine Manager:: It never (and I mean never) accepted a shutdown command. It would just ignore my command and keep chugging away. It got to the point I just used the force quit command. Never did find out why.

Getting Rid of a VM

Eventually, you'll be up a catalog of VMs you no longer want to tinker with.

Fortunately, getting rid of them in either VMS is no problem. Click the box in Boxes (or the three side dots if you opt for the list) and click delete.

In Virtual Machine Manager, right-click the heading and do the same.

Voila! Gone.

The End Analysis

For my testing I tried SparkyLinux GameOver Edition (3.7 GB). Chock full of games and assorted gizmos, it can be resource heavy.

I also tried Deepin OS, 5.4 GB and it, too, is resource heavy and a bear to load, even in live versions. Allegedly, it also has a lot of outgoing traffic (some claim it's China, the creator, spying on users).

Virtual Machine Manager loaded both up quickly, but could have been quicker if it wasn't for the elongated question and answer period beforehand.

Boxes was, essentially, a two or three mouse click affair until I was up and running with either OS. Definitely quicker and without the fuss.

However, Deepin killed both in short order. Virtual Machine Manager choked while picking a desktop design (required upon loading for the first time, or every instance in live versions), and Boxes gave up the fight after I attempted to open the app store.

SparkyLinux, even heavy with games, just chugged along in either one.

Which should you choose?

My advice is to try both and keep one. In my case, I chose Boxes only because it's so bog simple it's really hard to screw it up.

RETURN TO CONTENTS

Letters

Compiled by Ronnie Tucker

letters@fullcirclemagazine.org

RETURN TO CONTENTS

Q&A

Compiled by Erik The Unready

If you have a Linux question, email it to: misc@fullcirclemagazine.org, and Erik will answer them in a future issue. Please include as much information as you can about your query.

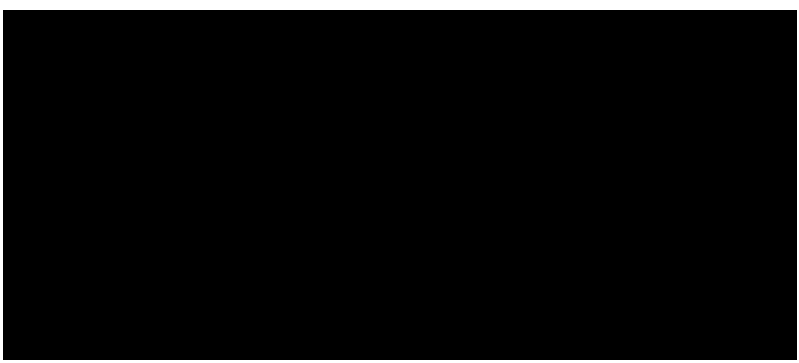

To understand the power of free software versus proprietary software, we can simply look at the final confirmation about listening in on cellphones. Facebook worms its way into every phone nowadays, to the point where you cannot remove it. Why? So they can get in your business. After confirming (what everyone knew anyway) that they listen to everything you say "for advertising purposes" (like that makes it any better!!!), I asked a cellphone user about it. His wife had this to say: "oh they can, I have nothing to hide", but I noticed she closes the bathroom door and the toilet door. His reply was a little more sensible: "but what can I do"? And it is true, with mobile phones we have no other choice, they have you by the short-n-curlies and abuse your trust. The problem is that the people who are supposed to protect you are in on the abuse. The

government or EU or someone needs to make laws and fine the crap out of Google and Huawei and Samsung and Meta and whatnot, but it is too juicy mining everyone's data, to put anyone who opposes the government in jail, instead of the criminals, while Big Tech gets away with Murder... If your phone OS was open source, you could block all of the nonsense yourself. I will get a cellphone one day when that happens and not before.

Q I switched to KDE neon as my hardware is getting on in years and Snap adds another 30s to my boot time. Why I'm asking, is because whenever I log on, my browser opens and I have to wait again. I can't seem to turn the autostarting off. Autostart has no entries. I even clicked on the "help", but nothing to solve my issue.

> A I think what you are looking for is not Autostart, but session restore. If you go to System settings -> Session, you should see "Session Restore", change that to "Start with an empty session" and click apply at the bottom-right.

Q I'm trying to virtualize Ubuntu 24.04 inside Ubuntu 24.04. I have the CPUs and memory for it, but every time I try to set it up in Virtualbox, I get: Failed to set the global VirtualBox extra data for key GUI/RecentFolderCD to value. Everything is updated, so I don't know what else to do. I have read that it is due to the folder naming, but I have my folder as "image" no apostrophes or anything: /home/dave/image I have even tried the latest version of Virtualbox not in the repositories, straight from the website, but it does the same thing.

> A I cannot tell you why, but I recall having that issue with backbox8.1, and not with backbox9. The smart thing will probably be to tell you to try a different hypervisor, but we can get a workaround. When you create the virtual machine, do not choose an image. Next-next until the VM is created. Now when you set up everything, choose the Ubuntu image in the storage options under the cd/dvd. It should launch after that.

Q I moved from Ubuntu Mate to Ubuntu Gnome this year, just to see how it would go. I usually do some light cosmetic changes, so I followed a few tutorials. What is the difference between ‘extensions’ and ‘extension manager’ ? Besides the fact that one's icon is green and the other blue?

> A There used to be bigger differences, now the only one is that “extension manager” allows you to download extensions from inside the application.

Q I upgraded to Ubuntu 24.4 now that it is stable, and I still run 22.4 in oracle virtualbox. Thing is, virtualbox only goes up to 23.10 if I choose my distro, and I want to install 24.4. If I choose 23.10 in the list and install 24.4, it does not complete, but fails with an error, “unsupported hypervisor”. So what should I use as a supported hypervisor for Ubuntu 22.4?

> A Usually that error is display settings related. Yes, I know, helpful error! Instead of choosing the default VMSVGA, choose VBOXSVGA and ignore the warning. It should load after that change.

Q I want to make my Ubuntu 24.04 install the same as my Ubuntu 20.04 install. I backed up the drive and installed a fresh copy of Ubuntu. Is there some way I can check out all the apps I had installed and install it all again as well as the repos, now that it is not installed?

> A If you are like me and do almost everything from the terminal, I suggest grabbing your “.bash_history” file from your backup and see what commands you used to set up your repositories and what you installed. Just remember to not just add repositories before checking their updates. Usually it should just go from Jammy to Noble, but check that!

Q I need help - like yesterday. I have a fresh install of Xubuntu and with Firefox I have 3 addons, privacy badger, ublock, and dark reader. Everything is fine, except I work at night and I use github a lot. All sites respect dark reader, except Github; it goes dark then suddenly white again, blasting my eyeballs. I’m using

Brave until it is sorted, but I have updated like 3x already, and it is still the same. I have even tried another dark addon, but same result. I'm not sure if others have this issue too, maybe it's just Github and Firefox?

> A You are not giving me much to go on, but it is safe to say your issue is with FireFox. Since you tried another dark reader (did you uninstall the old one first?), I'd say trash FF and get it again. My advice is to start with the plug-in, then the browser. It has nothing to do with Ubuntu.
>
> **EDIT** – I think I found it, I installed 'Dark reader' and clicked around randomly, but for anyone in the same boat, if you click on the name of the website, in this case, github.com inside of Dark reader, you can override the dark theme. So you may have clicked that by accident.

Q I have installed Ubuntu on my laptop from college and the i3 4GB beast with its 5400rpm drive barely runs Kubuntu. I wanted to use it for movies and series playing on my TV via HDMI. Thing is that no matter what player I use, celluloid, MPV or VLC, there are like micro-stutters. I thought it may be an update thing, so I connected it to the wireless and did an update, but it did nothing for the stutters and audio artifacts. Is it time for it to go to the landfill?

> A When you say you connected it to the network to update, I'm going to assume it was not connected and maybe also not during the install. That tells me that the codecs need to be installed. Try: sudo apt install ubuntu-restricted-extras

Q When I type lsb_release on Ubuntu 24.04, I get nothing back, eg. werner@amdpc:~$ lsb_release -v No LSB modules are available. Is this because I opted for the minimal install or do I need to use another command now?

> A Interesting, I checked for you and I also get the error, but I get information back: edd@gift:~$ lsb_release -cri

No LSB modules are available. Distributor ID: Ubuntu Release: 24.04 Codename: noble The man page says it all. Read the description paragraph. I am not sure *when it changed (but it feels like someone else also asked this recently).

Q I am using Ubuntu Gnome 24.04 on my laptop as the Mate version was too different from Mint that I was used to. I love Mint, but it is too far behind, so I opted for the flagship. One thing that I'm unsure of, is this pop-up I keep getting telling me there is a firmware update. I'm almost sure these things were taken care of by Mint during the OS update. I have updated at least 10x since the message appeared and it is not going away. Do I click on it? It now lives in the info panel of the date and time.

A Nothing is stopping you from not installing it if you do not want to? You can go to your manufacturer's website and see what the changes were and if you want them? I cannot tell you to install or not to install, that decision is yours to make. Look, nine times out of ten, there are fixes that allow for connecting new hardware and/or security fixes and it is not a bad thing.

Q Since I still live in the early 2000s, I have a PC not a laptop. I recently went and picked up bluetooth headphones from Takealot. I realised my PC does not have Bluetooth built in, so I grabbed an old Billionton Bluetooth class 1 dongle, from the early 2000s I had lying about. It works fine on Ubuntu 24.04, but I'm worried that it is too old and I will be hacked, should I?

A While there is always a risk of getting hacked, old hardware's flaws are usually more well known. Bluetooth effective distance is usually short, unlike WiFi, meaning that it should not really leave your house. What you could do is unplug the dongle when not in use? (a stitch in time and all that). You could also check if there is updated firmware for it? With that, I'd like to say keep your Ubuntu up-to-date and turn off your PC and not leave it on.

Q Hi there, I have Xubuntu 24.04 installed, quite recently. On my previous install of 20.04, I was able to set the brightness down quite far, but in 24.04, it seems as if it requires me to go grab the slider and do it by hand afterwards. Even then, it is not as dim as it was in 20.04. I work in low light, only the bit that comes in naturally, so the screen brightness on my laptop needs to be the absolute minimum. Is there another way for me to get it dimmer?

A There seems to be a project on github that does what you are asking, but I'm wondering, due to your first statement, if the amount of steps in your brightness slider is maybe set too low? If you click on the battery, below the slider there is a settings toggle; if you go there, just below the middle, you should see "brightness step count". If it is below twenty, I suggest setting it higher, the more steps there are, the lower it will go (10 steps = 0%, 10% and so on, where 20 steps = 0%, 5%, 10%, etc).

***Erik** has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.*

RETURN TO CONTENTS

Ubuntu Games: Fates of Ort

Written by Erik

Website: *https://www.fatesofort.com/*

Price: $4.49 (Gog sale) and $7.99(Steam)

Blurb: *“Fates of Ort is a retro fantasy RPG focused on strategic action, in a land where time is frozen when you stand still. Cast powerful spells - but beware, they will cost you your life.“*

Time flies when you’re having fun and it seems that Fates of Ort has gone from pre-release to post-release since we last had a look at it. I totally missed the 1.0, and I have 1.4.1 here. I am happy to report that the game has improved since then. There is still an issue from the pre-release, where enemies spawn every time you leave a screen and enter it again, but with the rest of the mechanics now working so smoothly, one could easily forgive that issue.

If you missed our first look at the game, Fates of Ort is an adventure game where you better be prepared to drop hours into it. It is *huge. It is truly open-world; usually, if you can see it, you can go there. (That does not mean you can just walk in a straight line, there are mountains and rivers and forests and swamps and deserts!). I will add a screenshot of the over-world here, but do not be fooled by it, each area has buildings you can enter, as well as underground locations not shown on the map, like caves and more! Only the very big towns are represented on the map.

There is the main quest, where you hunt your killer, and there are side quests everywhere.

Though the game is turn based, with an ‘I go, you go’ system, it mostly feels like real time, due to the way it was programmed; for instance, you can just walk, you don’t have to wait for other characters on the screen to take a step after you take a step, it just make the game feel more fluid. The game is 2D isometric, reminding me of the Ultima games, with a weird pixel art style for the characters. Again, they are faceless (what is it with artists and this impersonal style?), so I don’t really care for it. It makes the game feel a bit lifeless. With that said, navigating the world is easy and feels quite natural. I am a fan of isometric games, but I also know how easy it is to mess that up with characters going behind other objects. Here, that is kept to the bare minimum; the only time I lost my character was trying to go under the arch at the academy where you are killed, where it is blocked off at the other side, but you don’t know that until your character does not appear.

I seem to recall that the magic system was weird the last time I played it, but now, it is rather straightforward, I pick my spell and right-click; no more combining things to cast spells. The spells themselves are not just combat spells, there is a nice variety that you will use in other aspects of the game. Let’s talk about the magic system in the game quickly. There is no “mana” here, spells come off your life directly. This is a clever way to stop you from simply spamming the right-click or you

will drain your health. When casting the same spell over and over, it does start to increase in blood cost, so the game almost forces you to diversify your spell types. The health pool is generous though, so don't be disheartened. (He says as the Duke hits him for 900 damage and kills him instantly...).

The story also seems to have gotten the oil-change treatment, as small changes here and there feel like rather large changes, and things feel a bit more "together". There is still the annoying repeat issue, where you get the same stuff repeated when talking to characters, which, in my humble opinion, needs a clean up, but other than that, the game is a lot more fun to play than it was, what, like five years ago!?? Man, I cannot believe it has been that long since the last time I played this. Thus, a lot of things may feel a bit fuzzy and I may recall wrong, but bear with me.

Because the world does not move unless you do, you can take an action, like swinging your sword, to make time pass. This mechanism can be used to your advantage if you are smart, letting you get hits off without riposte. It also comes into play with slower spells, and you can lead enemies or influence their path. Later on, when you get the big boy spells, like meteor swarm, you need to take time into account as the spell needs a long time for the meteors to fall from the sky.

I have to mention the music, though it is not ear-gasmic; it does serve its purpose, and it does not play the same tune all the time ad nauseam; there are breaks and you can even set the

length of these. It leans into the retro aesthetic, sounding like music from early adventure games, maybe something like Chrono Trigger? So even if it is not farmer quality, (you know, outstanding in his field; yes, I know, exit stage left), the quality of the music is still top notch and very soothing, and it feels like there are a lot of tracks, so I cannot complain. (There are no OST or music files in the game folder, so I cannot say for sure how many tracks there are).

I think that it is the open world and the freedom you have to go anywhere that attracts me to this game, even though the featureless characters put me off. I mean, everything in the world has been given detail, but my character is a lump of clay. I'm not even joking here about the freedom you have, you can go straight to the end game from the start, bypassing everything else in the game. Is it a good idea? No, but you have that option open to you. Though the game is 2D isometric, it also feels a lot like Skyrim, where you cannot go a hundred steps without encountering "something", be it a house or a cave or an enemy or a side quest.

While my thoughts on this game are all over the place, the game itself is rather organised. Something I do not recall from the early access was the indication above other characters' heads, like the purple blobs indicating that they have had too much of the grape bubblegum that is popping up everywhere and will be hostile.

The game is made in the defold engine, which inspired me to

download defold and install it.

***Erik** has been in IT for 30+ years. He has seen technology come and go. From repairing washing machine sized hard drives with multimeters and oscilloscopes, laying cable, to scaling 3G towers, he's done it.*

RETURN TO CONTENTS

Patrons

The current site was created thanks to **Lucas Westermann** (Mr. Command & Conquer) who took on the task of completely rebuilding the site, and scripts, from scratch, in his own time. The Patreon page is to help pay the domain and hosting fees. The yearly target was quickly reached thanks to those listed on this page. The money also helps with the new mailing list that I set up.

Several people have asked for a PayPal (single donation) option, so I 've added a button to the right side of the website.

A big thank you to all those who've used Patreon and the PayPal button. It's a HUGE help.

https://www.patreon.com/fullcirclemagazine

https://paypal.me/ronnietucker

https://donorbox.org/recurring-monthly-donation

Monthly Patrons

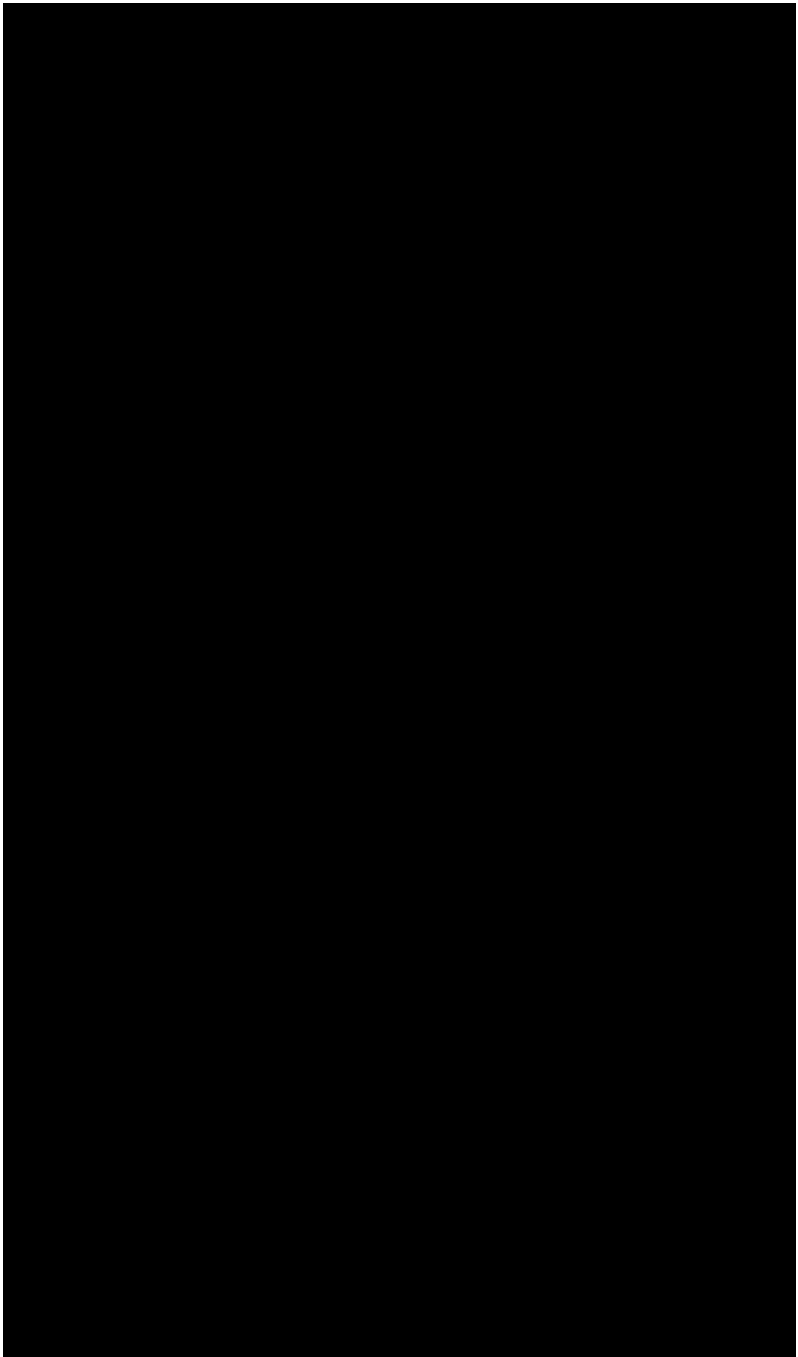

[redacted]

Single Donations

2024:

[redacted]

2025:

[redacted]

RETURN TO CONTENTS

How to Contribute

FULL CIRCLE NEEDS YOU!

*A magazine isn't a magazine without articles and Full Circle is no exception. We need your opinions, desktops, stories, how-to's, reviews, and anything else you want to tell your fellow *buntu users. Send your articles to:*
articles@fullcirclemagazine.org

We are always looking for new articles to include in Full Circle. For help and advice please see the ***Official Full Circle Style Guide*** *: http://bit.ly/fcmwriting*

FCM #216 Deadline: *Sunday 6th April 2025.*
Release: *Friday 25th April 2025.*

- *Send your comments or Linux experiences to: letters@fullcirclemagazine.org*
- *Hardware/software reviews should be sent to: reviews@fullcirclemagazine.org*
- *Questions for Q&A should go to: questions@fullcirclemagazine.org Desktop screens should be emailed to: misc@fullcirclemagazine.org*
- *… or you can visit our site via: fullcirclemagazine.org*

ronnie@fullcirclemagazine.org

admin@fullcirclemagazine.org

Getting Full Circle Magazine:

EPUB Format - Most editions of Full Circle have a link to the epub file on the downloads page. If you have any problems with the epub file, you can drop an email to: *mobile@fullcirclemagazine.org*

Issuu - You can read Full Circle online via Issuu: *http://issuu.com/fullcirclemagazine*. Please share and rate FCM as it helps to spread the word about FCM and Ubuntu Linux.

For the Full Circle Weekly News:

You can keep up to date with the Weekly News using the RSS feed: *http://fullcirclemagazine.org/feed/podcast*

Or, if you're out and about, you can get the Weekly News via Stitcher Radio (Android/iOS/

web): *http://www.stitcher.com/s?fid=85347&refid=stpr*
and via TuneIn at: *http://tunein.com/radio/Full-Circle-Weekly-News-p855064/*

FCM PATREON : https://www.patreon.com/fullcirclemagazine

RETURN TO CONTENTS